



US006650325B1

(12) **United States Patent**  
**Voorhies et al.**

(10) **Patent No.: US 6,650,325 B1**  
(45) **Date of Patent: Nov. 18, 2003**

(54) **METHOD, APPARATUS AND ARTICLE OF MANUFACTURE FOR BOUSTROPHEDONIC RASTERIZATION**

(75) Inventors: **Douglas A. Voorhies**, Menlo Park, CA (US); **Nicholas J. Foskett**, Mountain View, CA (US)

(73) Assignee: **NVIDIA Corporation**, Santa Clara, CA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/454,505**

(22) Filed: **Dec. 6, 1999**

(51) **Int. Cl.**<sup>7</sup> ..... **G06T 15/40**; G06T 15/50; G06T 17/00

(52) **U.S. Cl.** ..... **345/426**; 345/422; 345/428; 345/581; 345/586; 345/552; 345/561

(58) **Field of Search** ..... 345/418-422, 345/426, 428, 468, 469, 501, 505, 502, 506, 581, 582, 586-587, 600, 601, 610, 611, 612, 613, 619, 519, 520, 522, 530, 531, 545-549, 552, 556, 204, 561, 553, 615, 627, 440-441, 643, 469.1

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

|           |    |   |         |                 |       |         |
|-----------|----|---|---------|-----------------|-------|---------|
| 5,313,287 | A  | * | 5/1994  | Barton          | ..... | 358/458 |
| 5,598,517 | A  |   | 1/1997  | Watkins         | ..... | 395/141 |
| 5,694,143 | A  |   | 12/1997 | Fielder et al.  |       |         |
| 5,821,944 | A  |   | 10/1998 | Watkins         | ..... | 345/430 |
| 5,838,337 | A  |   | 11/1998 | Kimura et al.   |       |         |
| 5,977,997 | A  |   | 11/1999 | Vainsencher     |       |         |
| 6,000,027 | A  |   | 12/1999 | Pawate et al.   |       |         |
| 6,057,855 | A  | * | 5/2000  | Barkans         | ..... | 345/435 |
| 6,166,717 | A  | * | 12/2000 | Nakayama        | ..... | 345/611 |
| 6,198,488 | B1 | * | 3/2001  | Lindholm et al. | ..... | 345/426 |
| 6,236,408 | B1 |   | 5/2001  | Watkins         | ..... | 345/433 |
| 6,285,779 | B1 | * | 9/2001  | Lapidous et al. | ..... | 382/106 |

|           |    |   |         |                 |       |         |
|-----------|----|---|---------|-----------------|-------|---------|
| 6,288,722 | B1 | * | 9/2001  | Narayanaswami   | ..... | 345/619 |
| 6,342,888 | B1 | * | 1/2002  | Lindholm et al. | ..... | 345/426 |
| 6,353,439 | B1 | * | 3/2002  | Lindholm et al. | ..... | 345/561 |
| 6,417,851 | B1 | * | 7/2002  | Lindholm et al. | ..... | 345/426 |
| 6,417,858 | B1 |   | 7/2002  | Bosch et al.    | ..... | 345/522 |
| 6,452,595 | B1 | * | 9/2002  | Montrym et al.  | ..... | 345/426 |
| 6,462,737 | B2 | * | 10/2002 | Lindholm et al. | ..... | 345/426 |
| 6,476,807 | B1 | * | 11/2002 | Duluk et al.    | ..... | 345/421 |
| 6,504,542 | B1 | * | 1/2003  | Voorhies et al. | ..... | 345/441 |
| 6,515,671 | B1 | * | 2/2003  | Lindholm et al. | ..... | 345/506 |

**FOREIGN PATENT DOCUMENTS**

|    |           |    |         |
|----|-----------|----|---------|
| EP | 0609430   | A2 | 1/1996  |
| EP | 0690430   | A3 | 7/1996  |
| JP | 10-269-66 |    | 1/1998  |
| JP | 10-40679  |    | 2/1998  |
| WO | 98/28695  |    | 7/1998  |
| WO | 99/52040  |    | 10/1999 |

**OTHER PUBLICATIONS**

Marc Olano and Trey Greer; "Triangle Scan Conversion Using 2D Homogeneous Coordinates"; 1997, SIGGRAPH/Eurographics Workshop.

\* cited by examiner

*Primary Examiner*—Matthew C. Bella

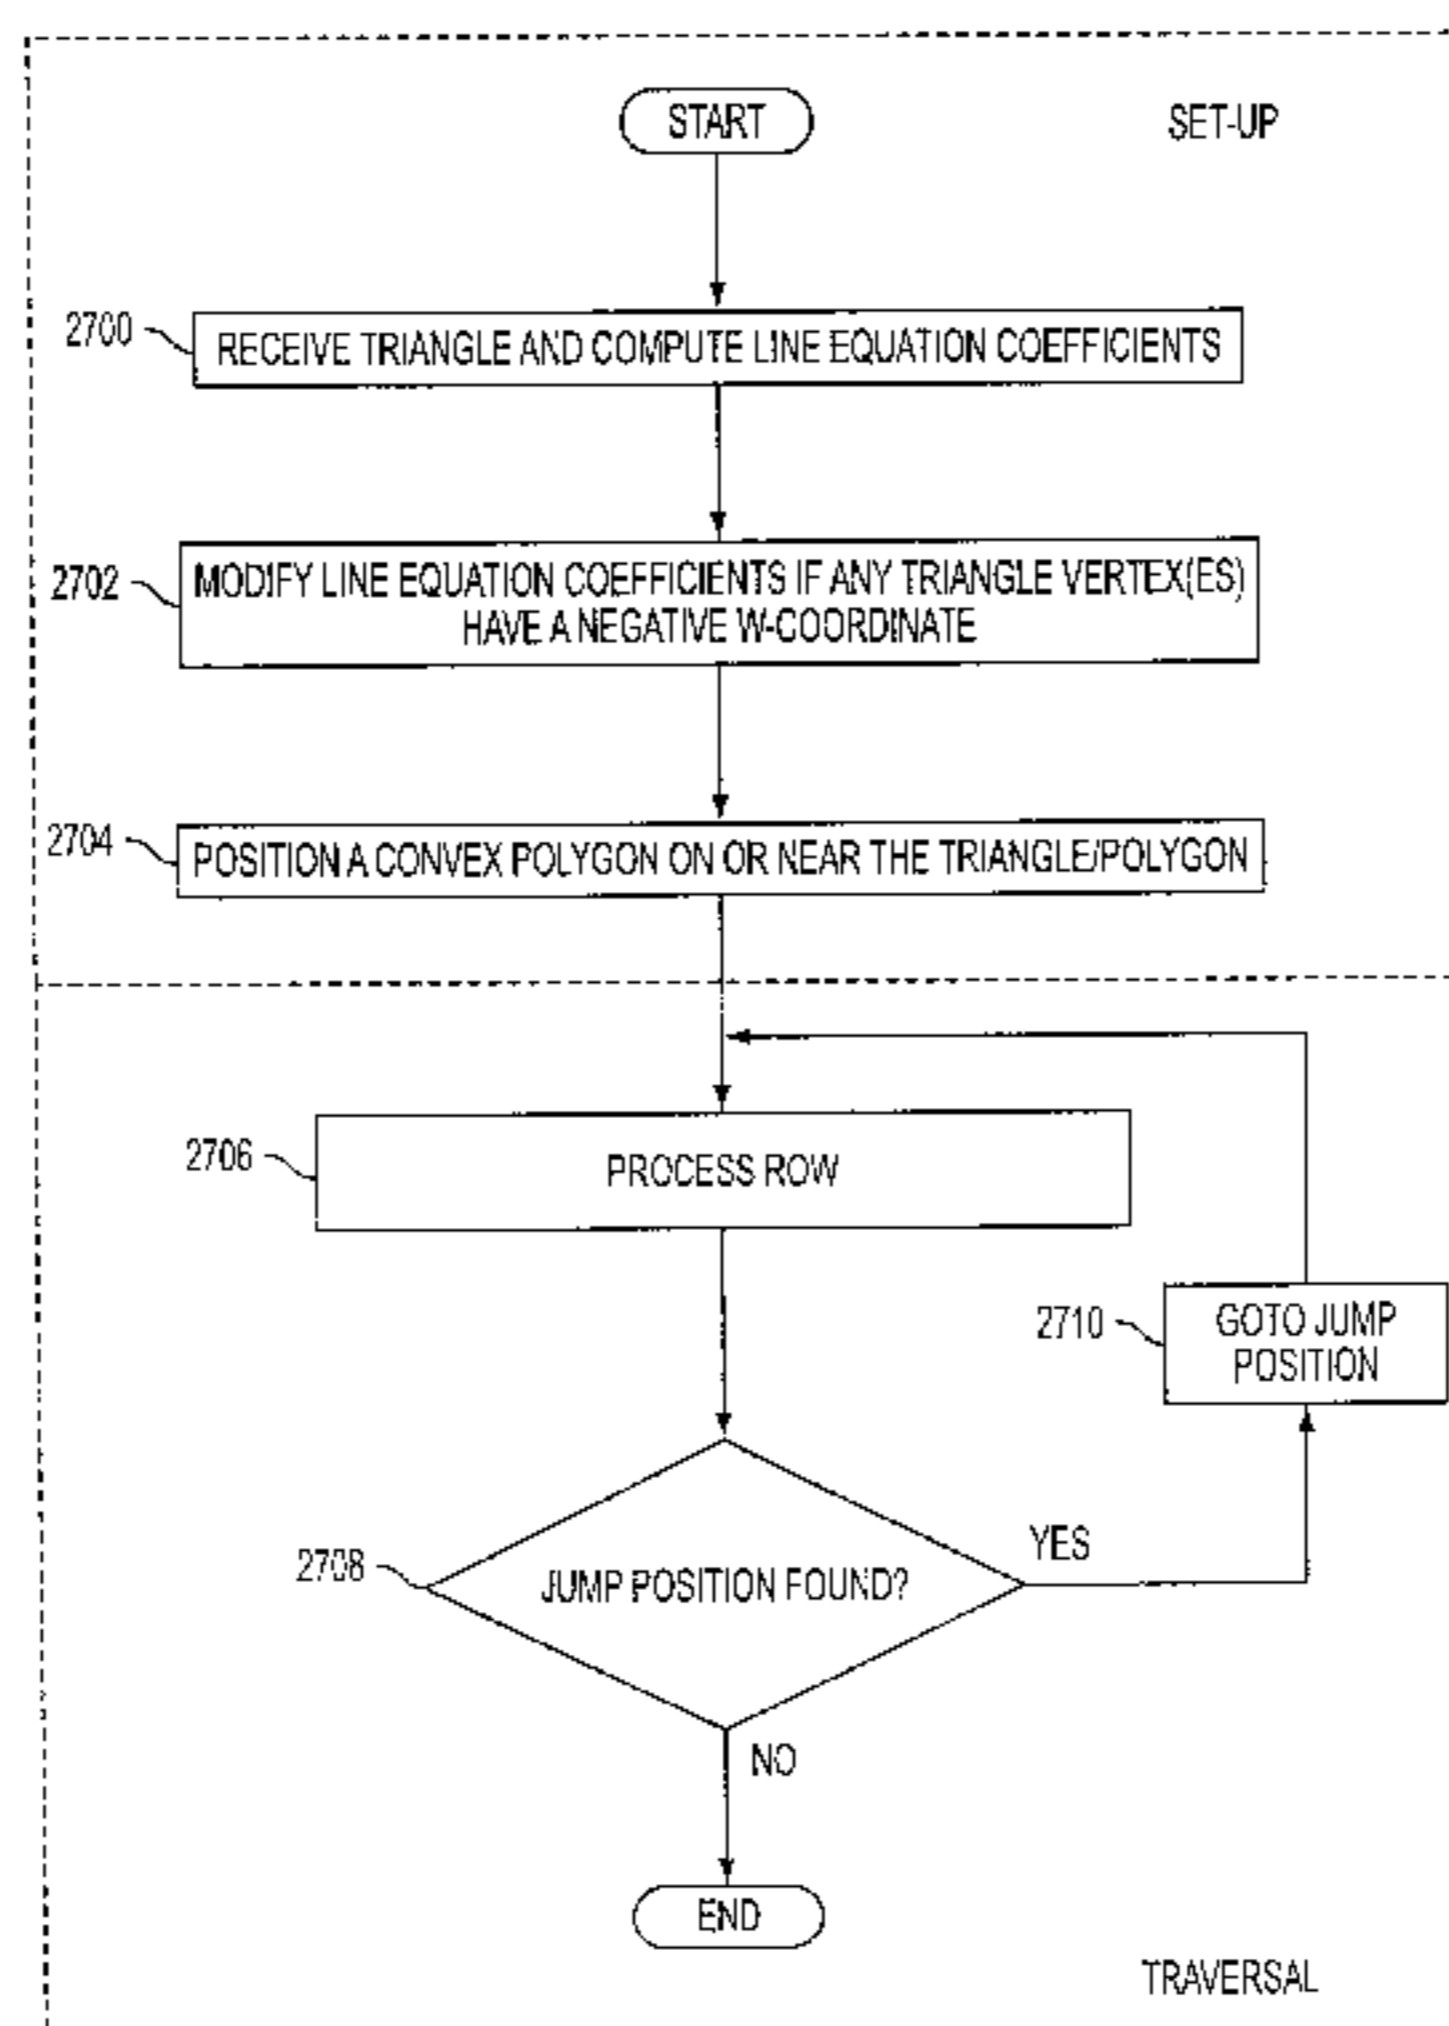
*Assistant Examiner*—Wesner Sajous

(74) *Attorney, Agent, or Firm*—Silicon Valley IP Group; Kevin J. Zilka

(57) **ABSTRACT**

A method, apparatus and article of manufacture are provided for performing rasterization using alternating sense point traversal. Upon receipt of a primitive, i.e. a triangle, a plurality of points are positioned on or near the primitive. Such points define an enclosed convex region and may be located at corners of the convex region. In operation, the points and convex region are moved in an alternating manner for the purpose of identifying an area in the primitive for rendering pixels therein. In particular, the points are moved in a boustrophedonic manner.

**35 Claims, 43 Drawing Sheets**






FIG. 1  
(PRIOR ART)




FIG. 1A  
(PRIOR ART)




FIG. 1B




FIG. 2

| COMMAND           | TRANSFORM STALL | LIGHTING STALL | DESCRIPTION   |
|-------------------|-----------------|----------------|---|
| FE2XF_CMD_NOP     |                 |                | NO OPERATION. CAN BE USED AS A SPACER BETWEEN COMMANDS.         |
| FE2XF_CMD_VERTEX  | READ            | READ           | VERTEX DATA.  |
| FE2XF_CMD_PASSTHR |                 |                | PASSTHROUGH. TRANSFORM AND LIGHTING PASS THE DATA THROUGH.      |
| FE2XF_CMD_RDVAB   |                 |                | READ THE VAB CONTENTS WHEN CONTEXT SWITCHING.                   |
| FE2XF_CMD_LDMODE  |                 |                | LOAD NEW MODE BITS.   |
| FE2XF_CMD_LDXFCTX | WRITE           |                | LOAD TRANSFORM CONTEXT MEMORY DATA.                             |
| FE2XF_CMD_RDXFCTX | READ            |                | READ TRANSFORM CONTEXT MEMORY DATA.                             |
| FE2XF_CMD_LDLTCTX |                 | WRITE          | LOAD LIGHTING CONTEXT MEMORY DATA.                              |
| FE2XF_CMD_RDLTCTX |                 | READ           | READ LIGHTING CONTEXT MEMORY DATA.                              |
| FE2XF_CMD_LDLTC0  |                 | WRITE          | LOAD LIGHTING CONTEXT0 MEMORY DATA.                             |
| FE2XF_CMD_RDLTC0  |                 | READ           | READ LIGHTING CONTEXT0 MEMORY DATA.                             |
| FE2XF_CMD_LDLTC1  |                 | WRITE          | LOAD LIGHTING CONTEXT1 MEMORY DATA.                             |
| FE2XF_CMD_RDLTC1  |                 | READ           | READ LIGHTING CONTEXT1 MEMORY DATA.                             |
| FE2XF_CMD_LDLTC2  |                 | WRITE          | LOAD LIGHTING CONTEXT2 MEMORY DATA.                             |
| FE2XF_CMD_RDLTC2  |                 | READ           | READ LIGHTING CONTEXT2 MEMORY DATA.                             |
| FE2XF_CMD_LTLTC3  |                 | WRITE          | LOAD LIGHTING CONTEXT3 MEMORY DATA.                             |
| FE2XF_CMD_RDLTC3  |                 | READ           | READ LIGHTING CONTEXT3 MEMORY DATA.                             |
| FE2XF_CMD_SYNC    | READ+WRITE      | READ+WRITE     | SIMILAR TO NOP, BUT IS NOT ALLOWED TO BE PROCESSED IN PARALLEL. |

FIG. 2A



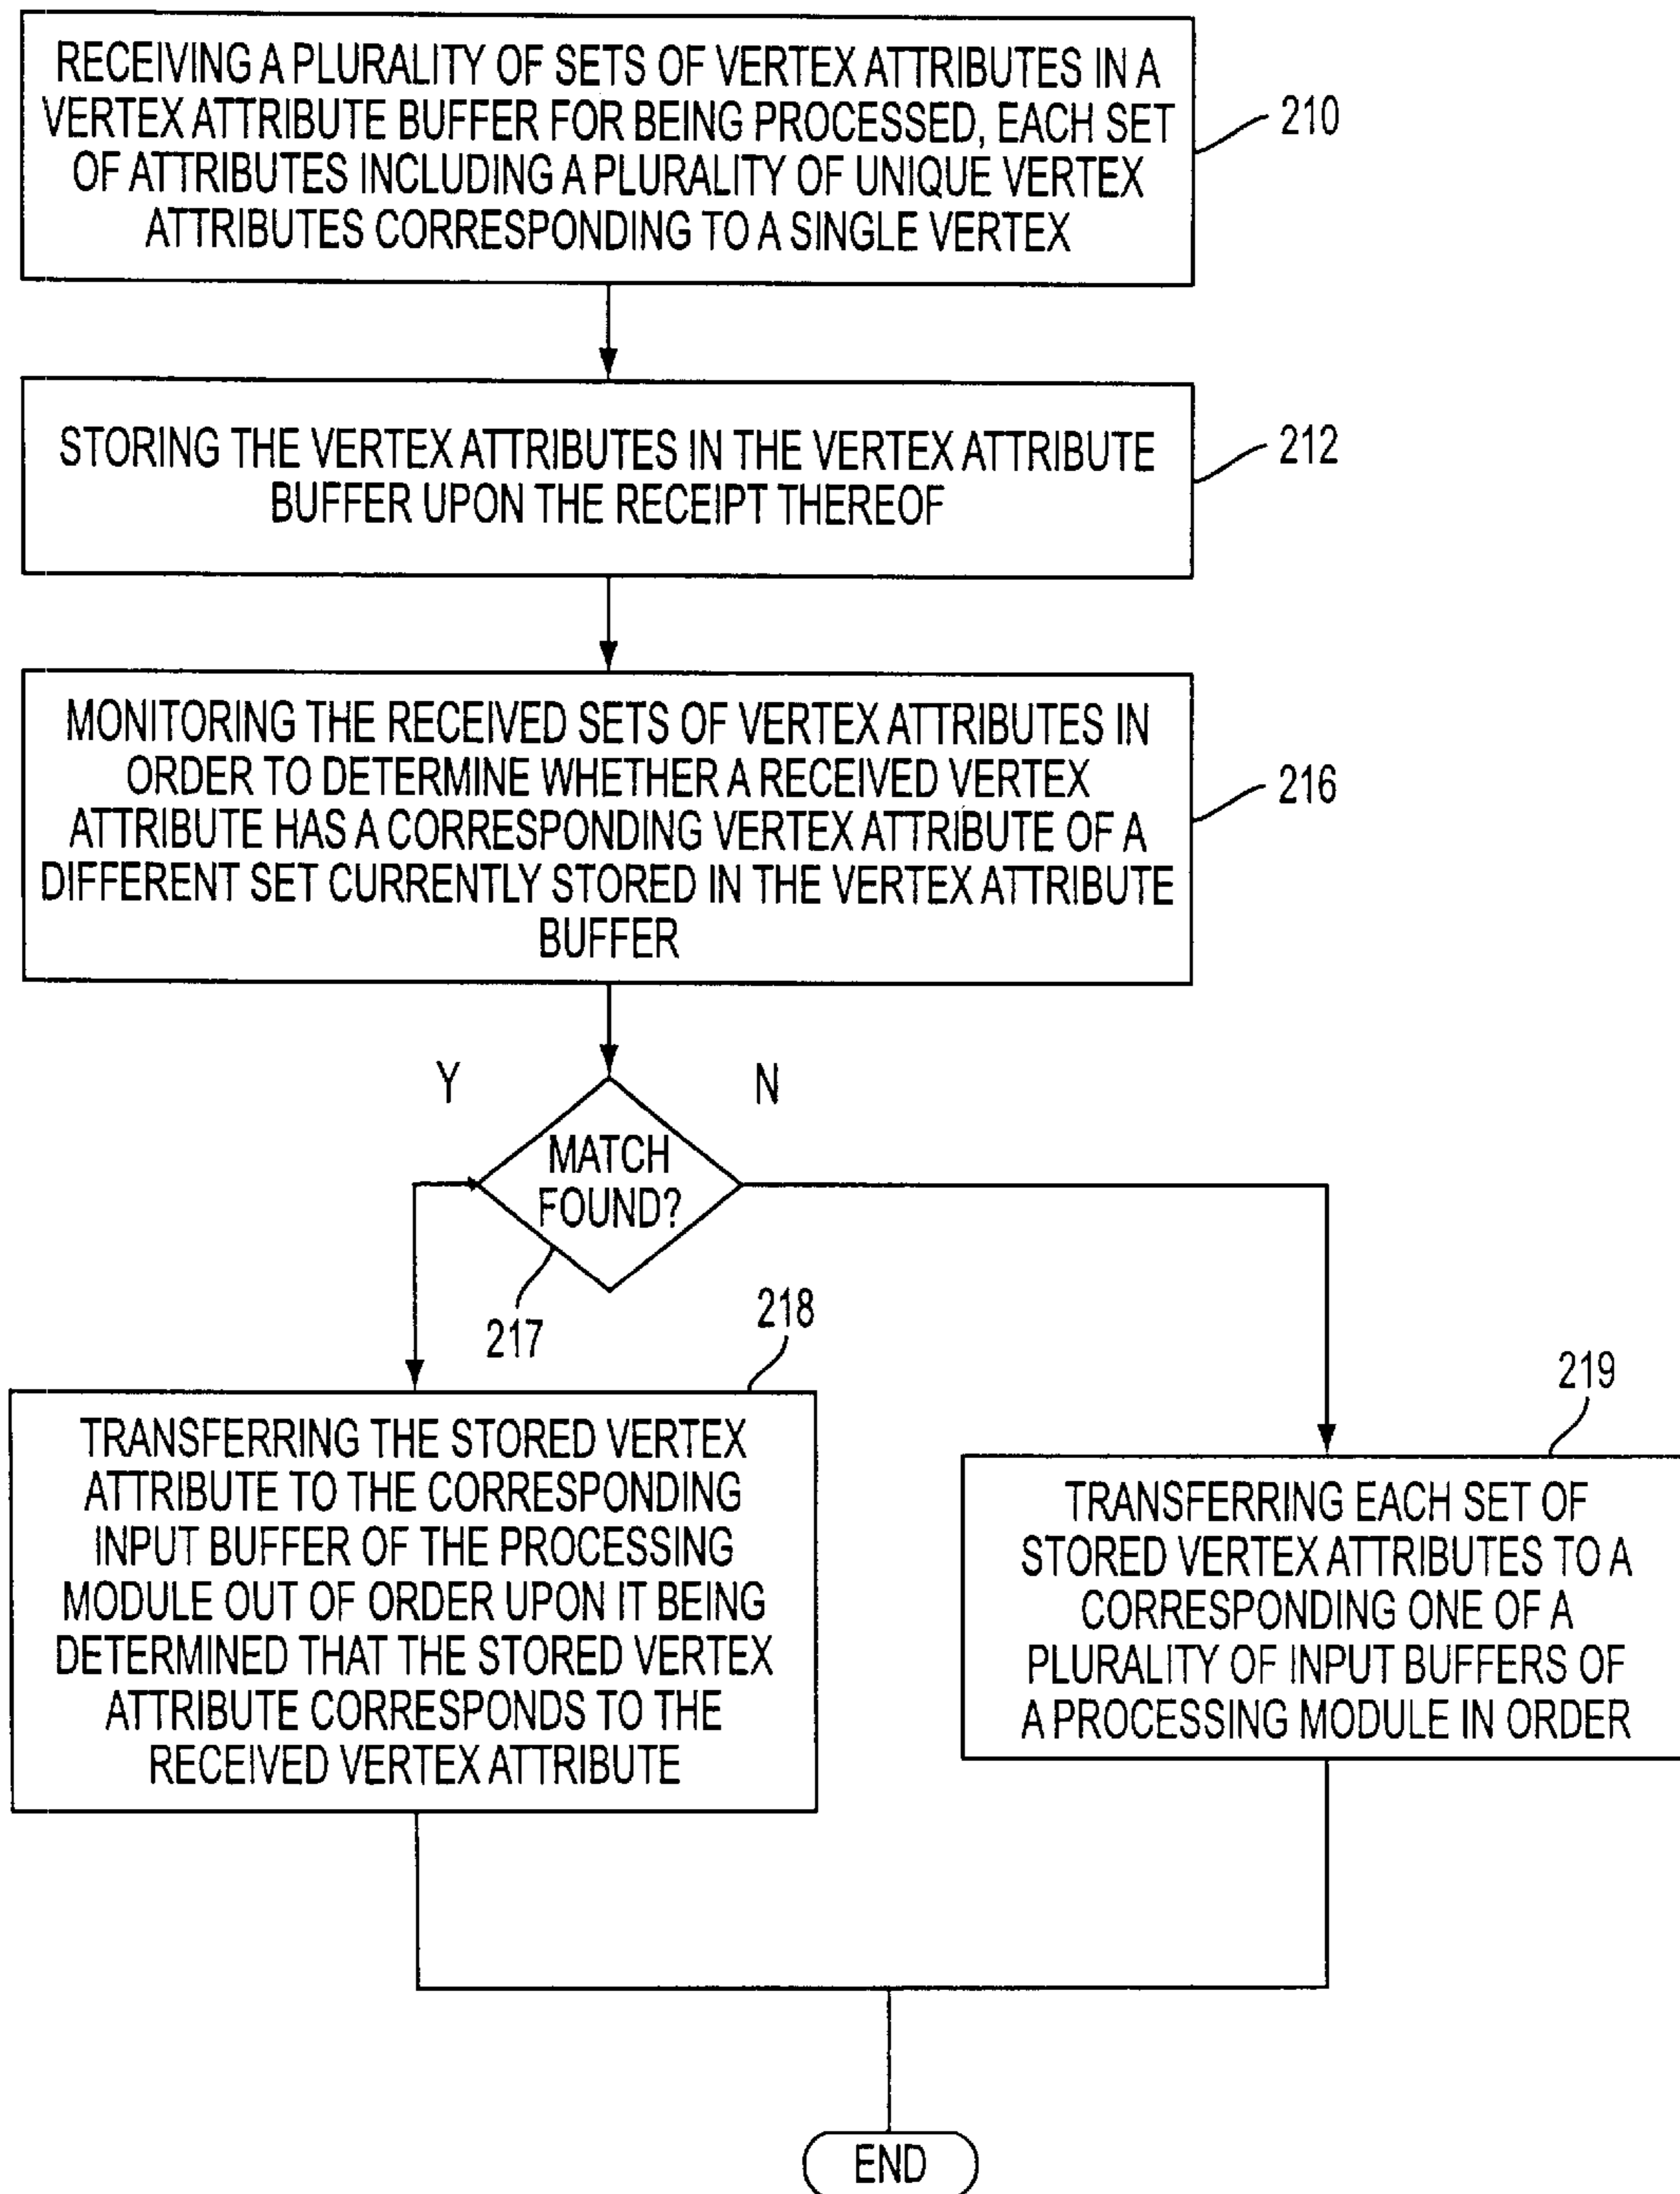


FIG. 2B

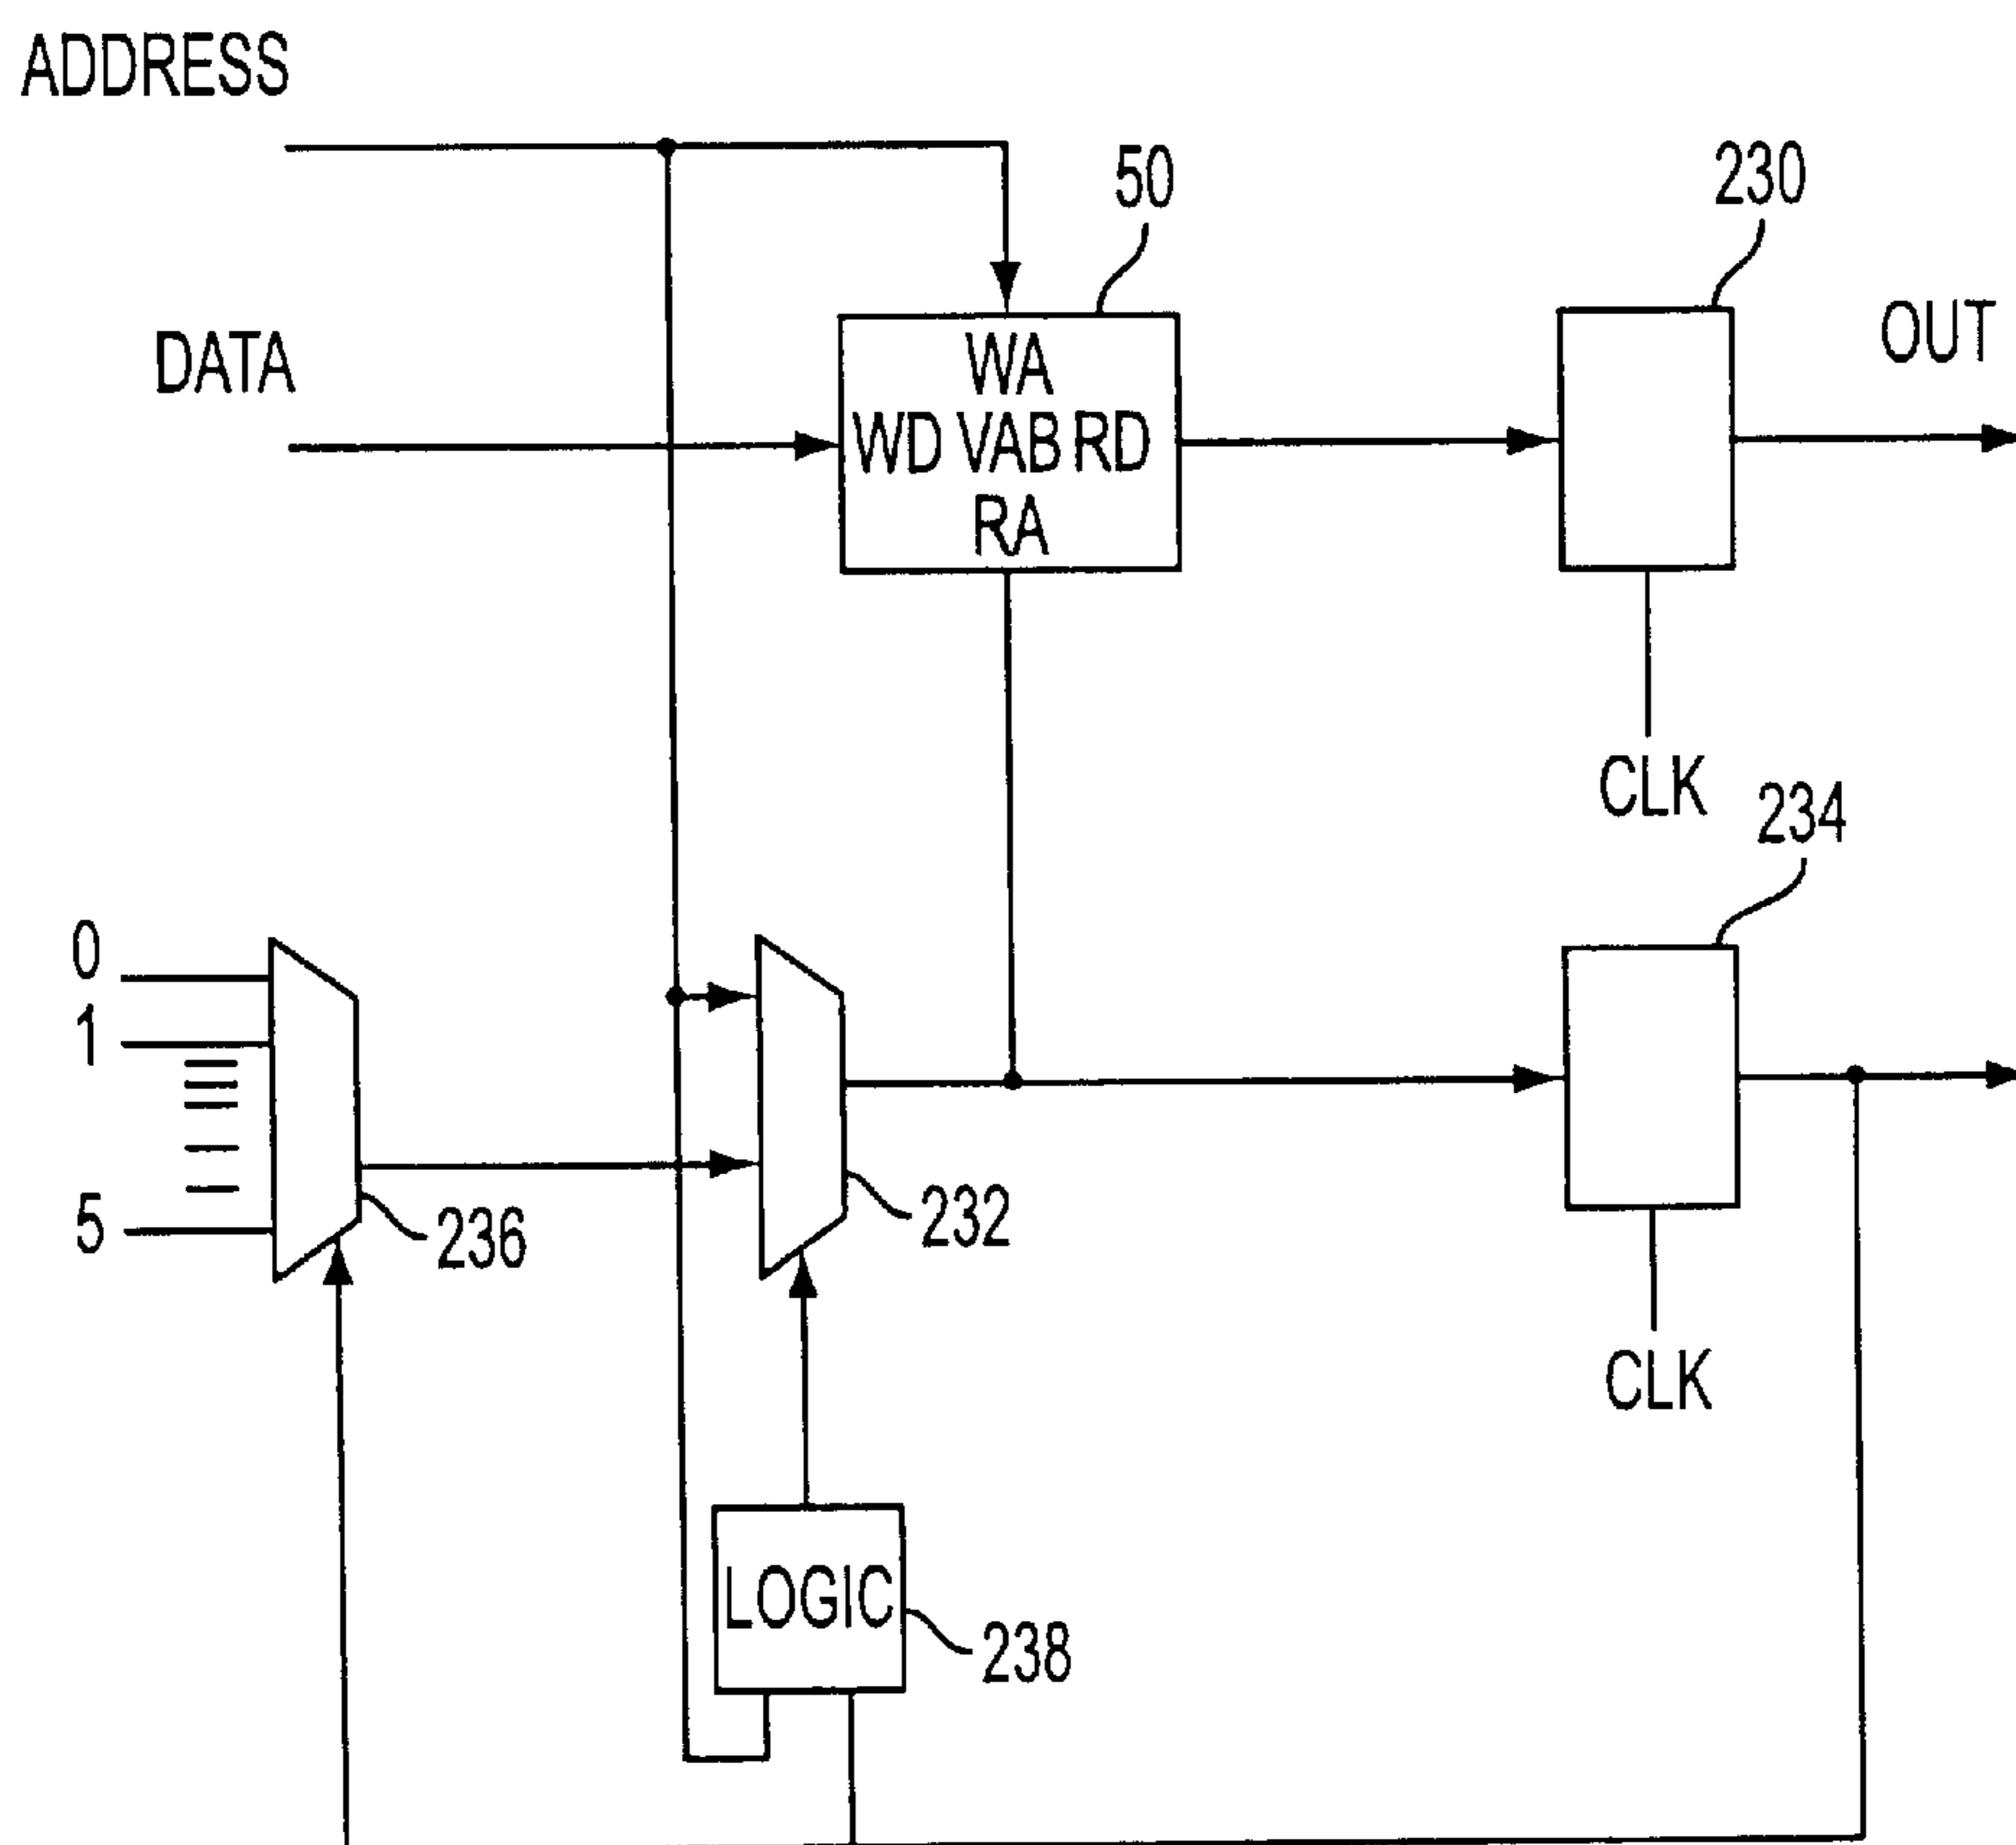


FIG. 2C

| MODE BIT | BITS | DESCRIPTION  |
|----------|------|--|
| T0       | 1    | TEXTURE 0 ENABLE   |
| TXF0     | 1    | TEXTURE 0 MATRIX TRANSFORM ENABLE  |
| TDV0     | 1    | TEXTURE 0 w DIVIDE ENABLE  |
| T0S      | 3    | TEXTURE 0 TEXGEN s CONTROL   |
| T0T      | 3    | TEXTURE 0 TEXGEN t CONTROL   |
| T0U      | 3    | TEXTURE 0 TEXGEN r CONTROL   |
| T0Q      | 2    | TEXTURE 0 TEXGEN q CONTROL   |
| T1       | 1    | TEXTURE 1 ENABLE   |
| TXF1     | 1    | TEXTURE 1 MATRIX TRANSFORM ENABLE  |
| TDV1     | 1    | TEXTURE 1 w DIVIDE ENABLE  |
| T1S      | 3    | TEXTURE 1 TEXGEN s CONTROL   |
| T1T      | 3    | TEXTURE 1 TEXGEN t CONTROL   |
| T1U      | 3    | TEXTURE 1 TEXGEN r CONTROL   |
| T1Q      | 2    | TEXTURE 1 TEXGEN q CONTROL   |
| ETY      | 1    | EYE TYPE INFINITE(0) OR LOCAL(1)   |
| LIT      | 1    | LIGHTING ENABLE  |
| NRM      | 1    | NORMAL NORMALIZE ENABLE  |
| FOG      | 1    | FOG ENABLE   |
| LIS      | 16   | LIGHT STATUS ( 8 LIGHTS BY 2 BITS EACH,<br>0:OFF,1:INFINITE,2:LOCAL,3:SPOTLIGHT) |
| FG       | 2    | FOGGEN CONTROL(0: OFF, 1:RADIAL, 2: PLANE)                                       |
| LAT      | 1    | LIGHT ATTENUATION CONTROL (0: INVERT, 1: NO INVERT)                              |
| C1I      | 1    | SPECULAR COLOR INPUT ENABLE  |
| C1O      | 1    | SPECULAR COLOR OUTPUT ENABLE   |
| CM       | 4    | COLOR MATERIAL CONTROL (1: EMISSIVE, 2: AMBIENT, 4:<br>DIFFUSE, 8:SPECULAR)      |
| PP       | 1    | POINT PARAMETER ENABLE   |
| SKIN     | 1    | SKINNING ENABLE  |
| VPAS     | 1    | VERTEX PASS ENABLE   |

FIG. 3




FIG. 4



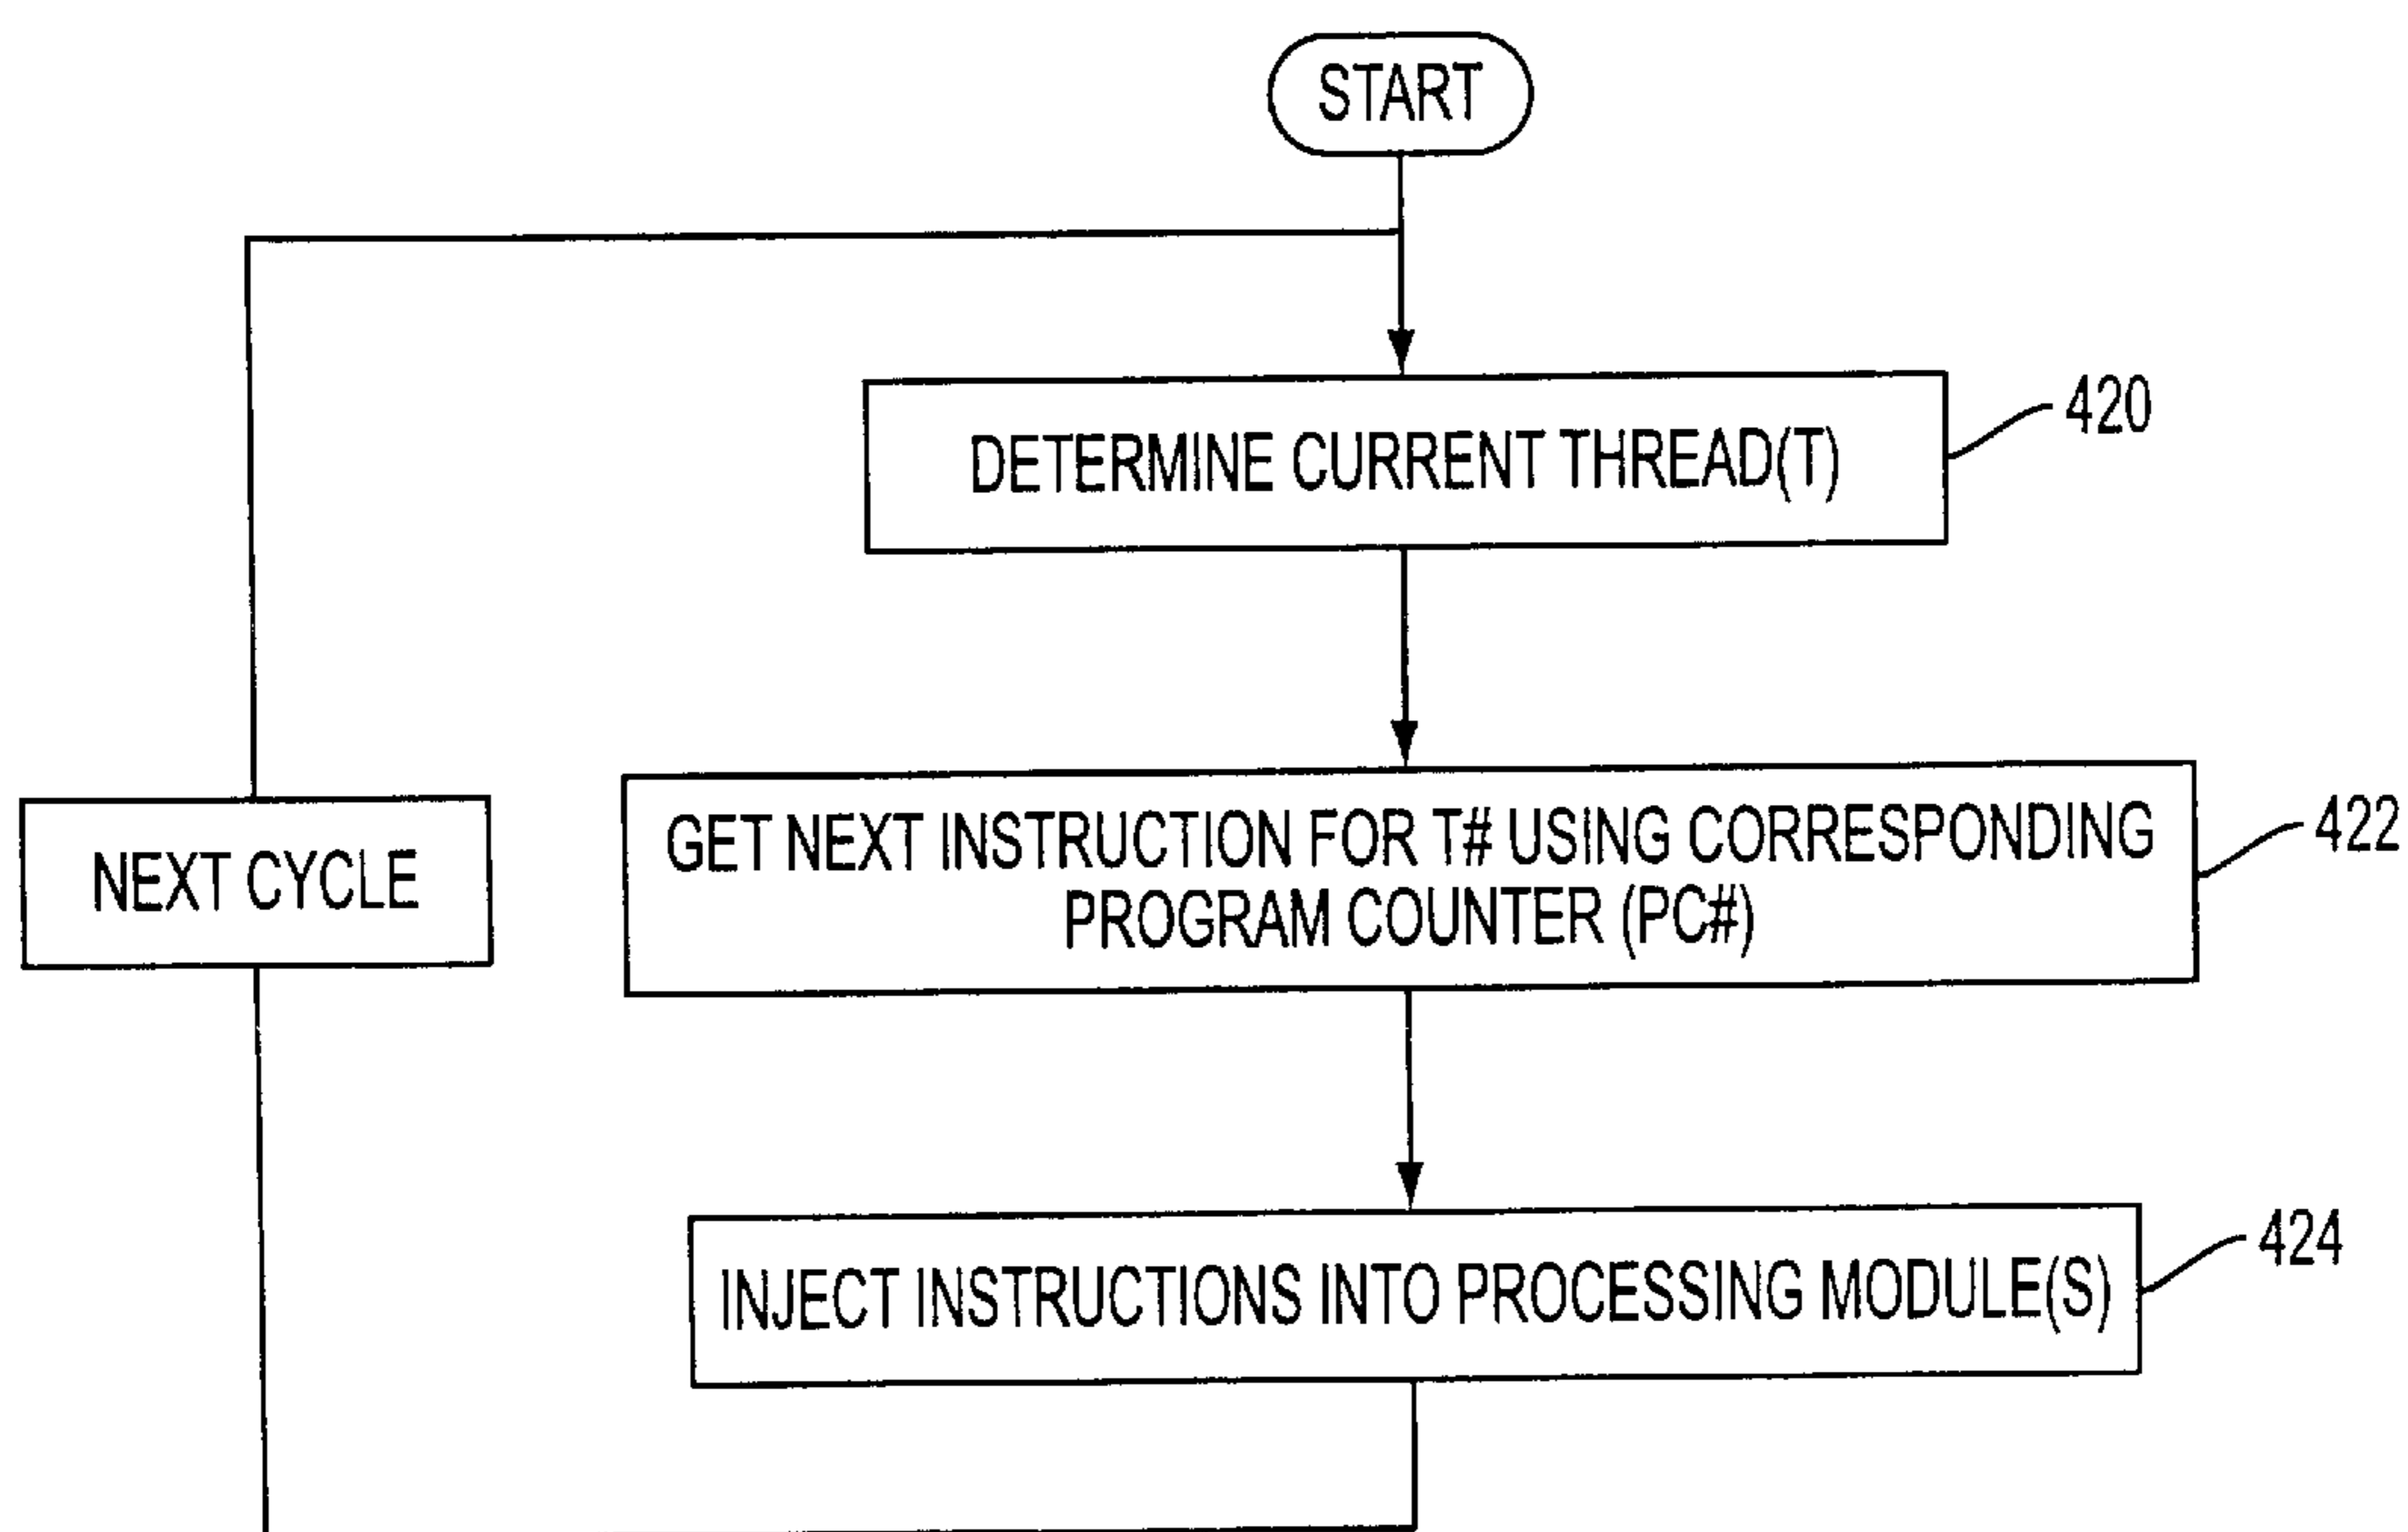


FIG. 4A




FIG. 4B




FIG. 5

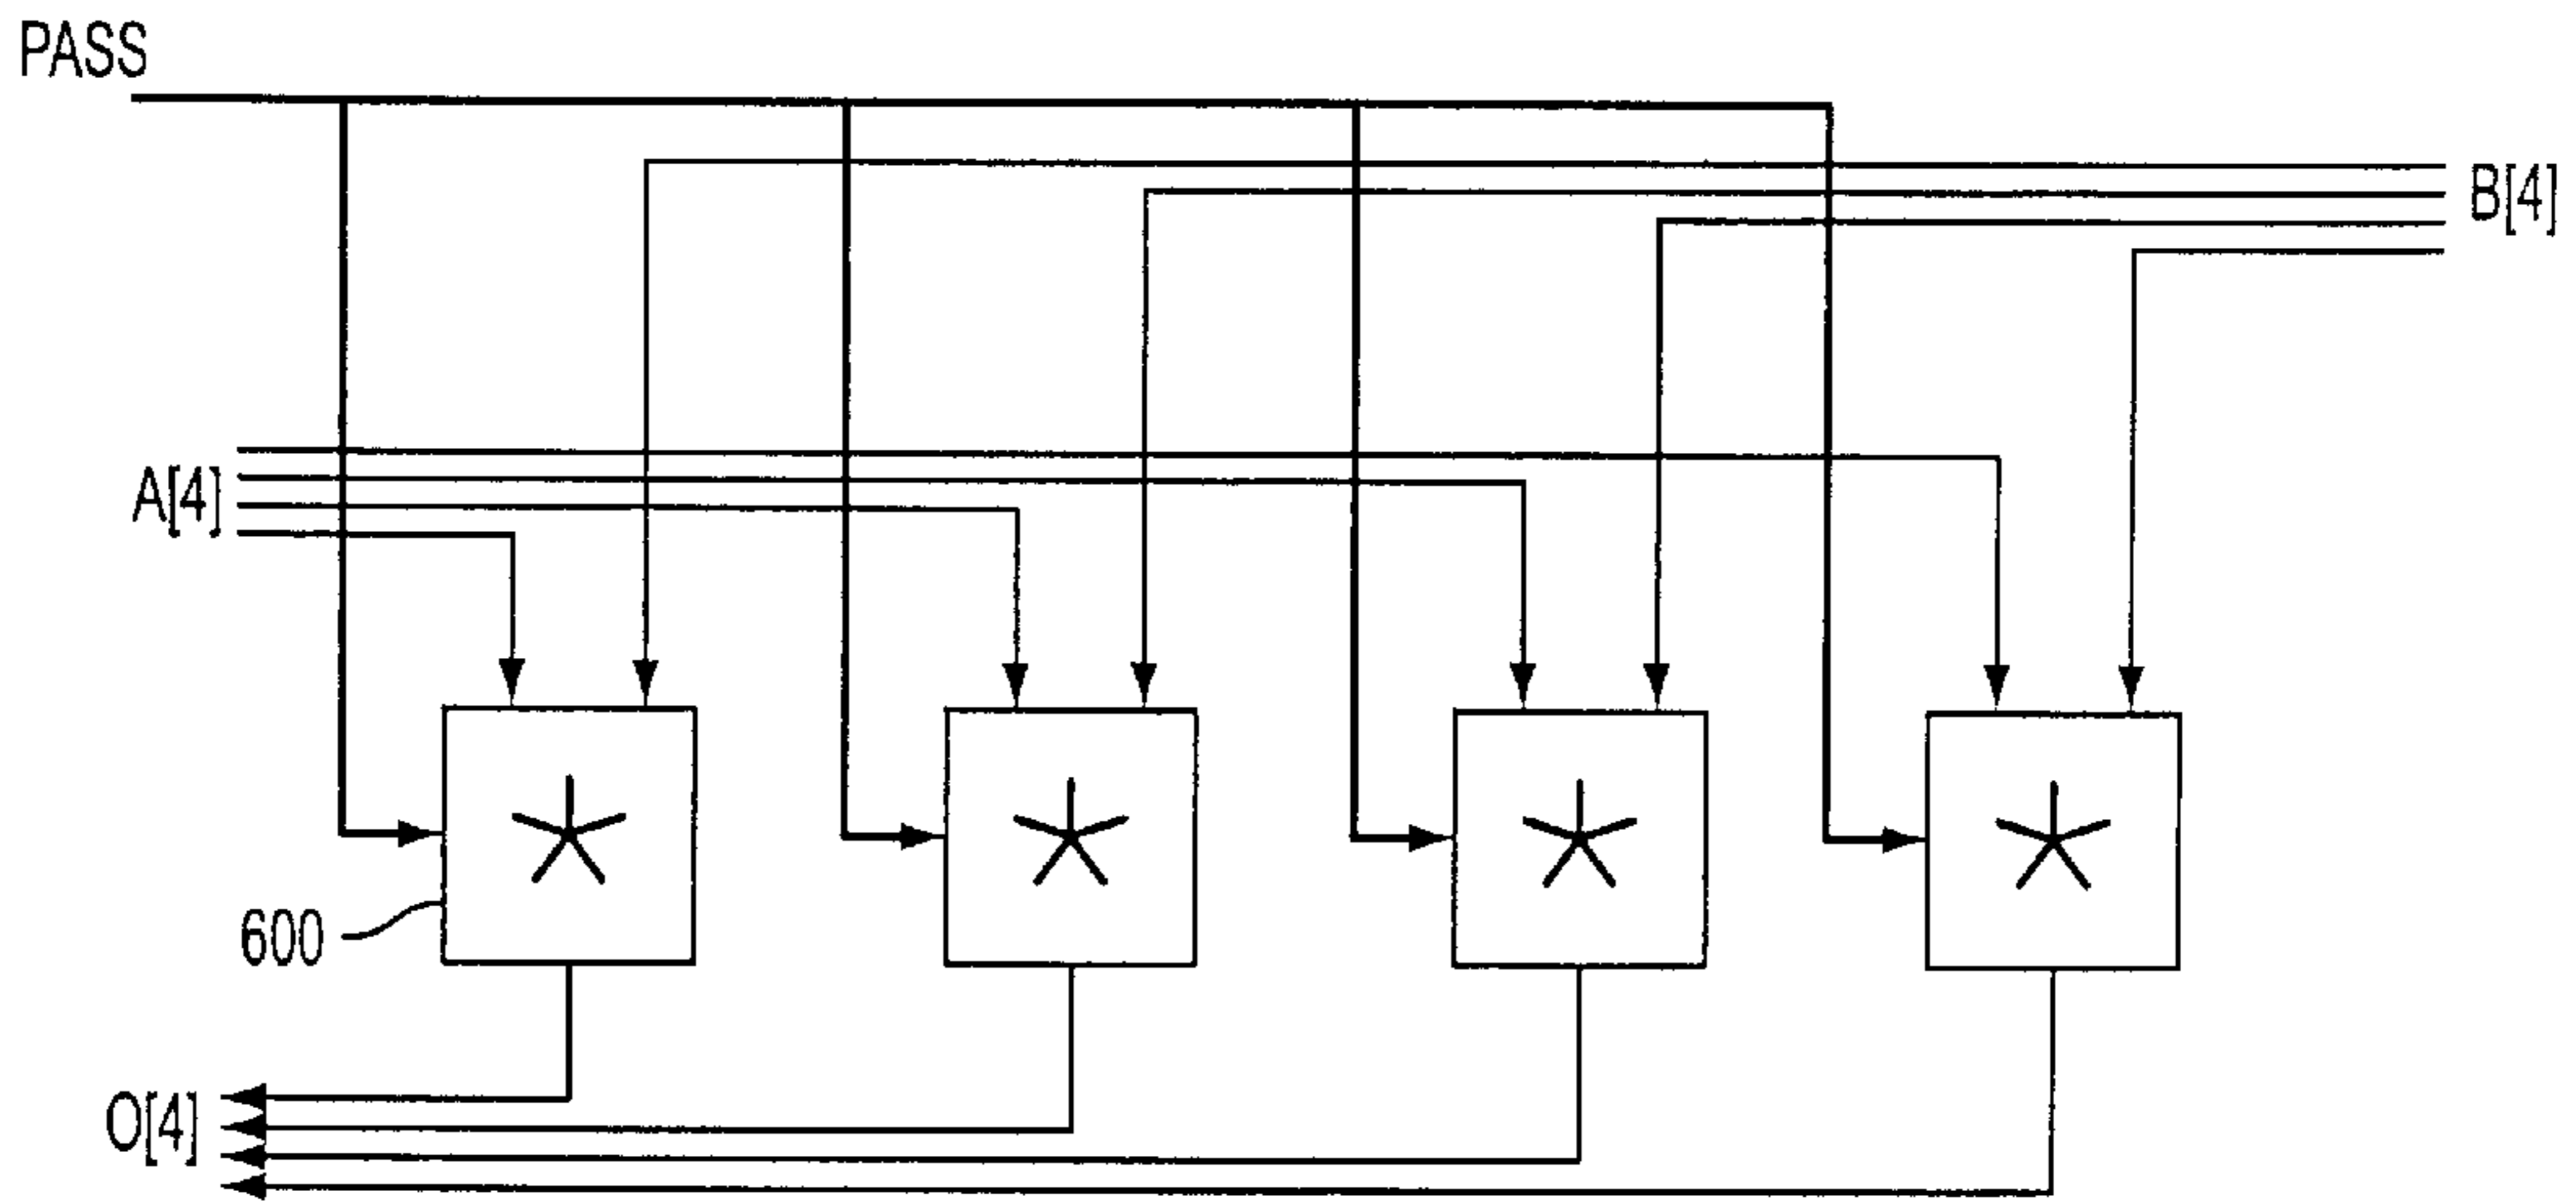


FIG. 6

500

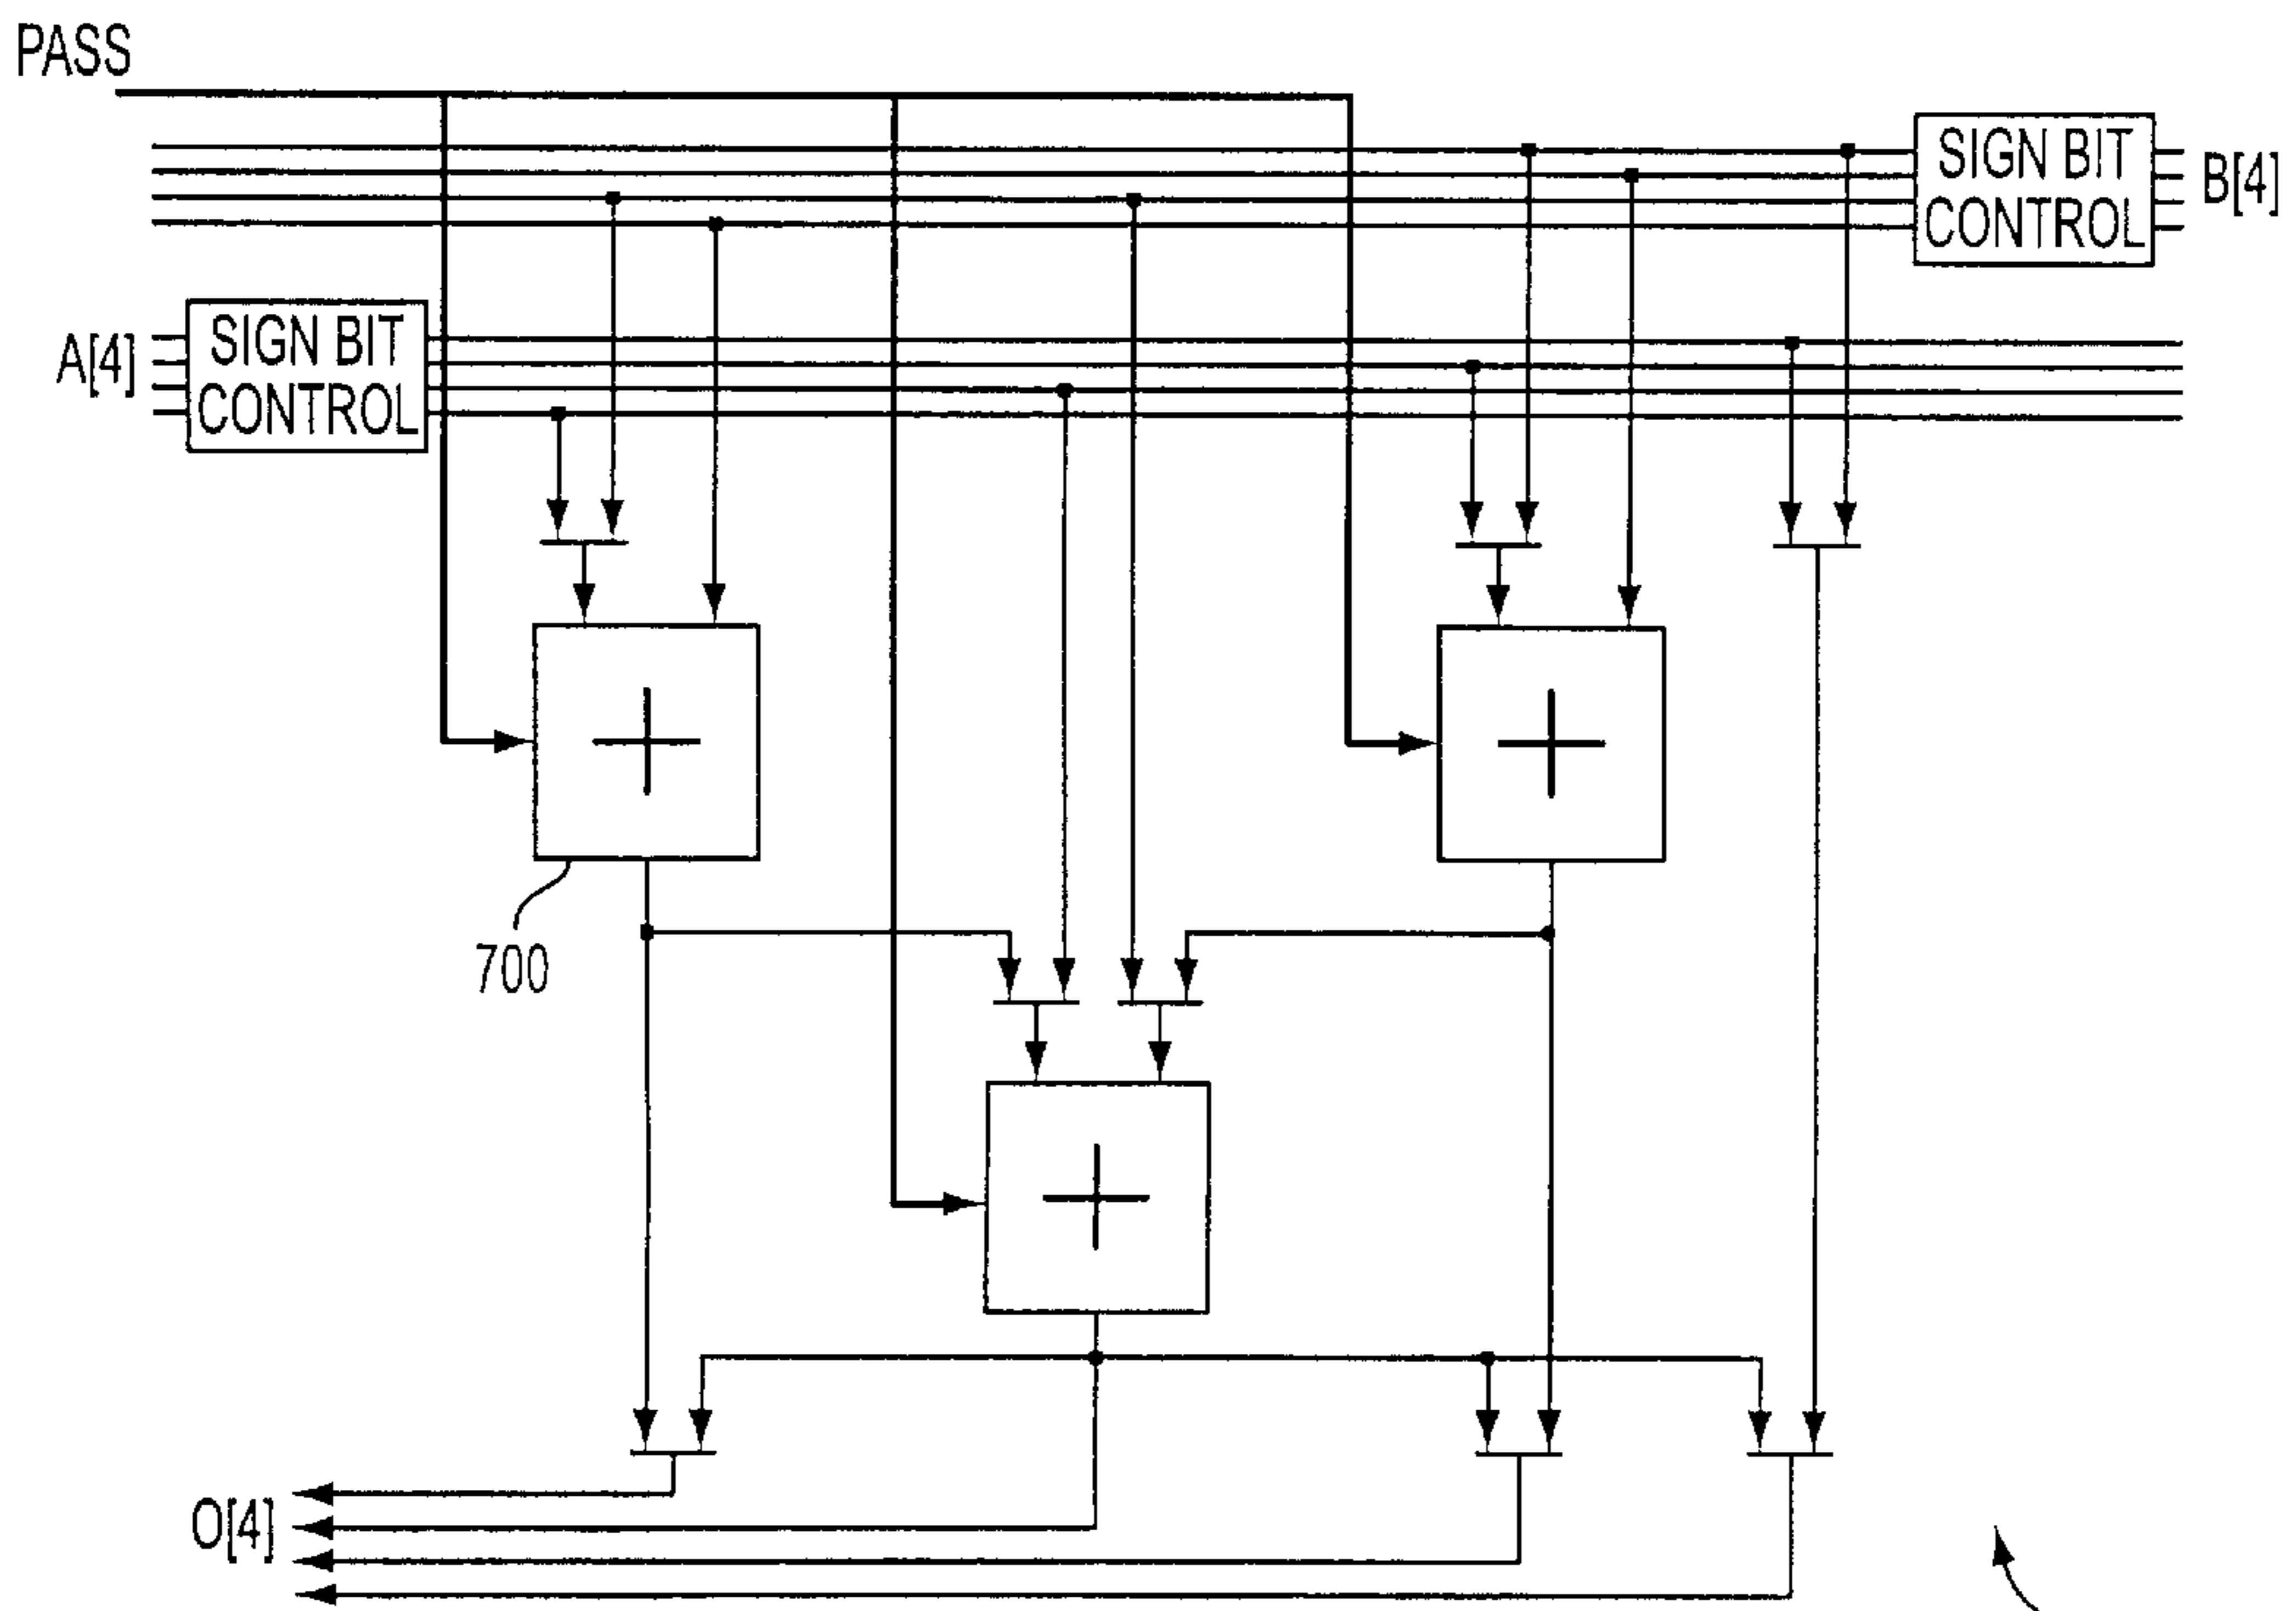


FIG. 7

504




FIG. 8




FIG. 9



| ADDRESS | TARGET          | ACTION                       | DESCRIPTION                |
|---------|-----------------|------------------------------|----------------------------|
| TPOS    | TBUFFER         | T[0] = ALU                   | POSITION                   |
| TT0     | TBUFFER         | T[3] = ALU                   | TEXTURE0                   |
| TT1     | TBUFFER         | T[4] = ALU                   | TEXTURE1                   |
| WEV     | WBUFFER,VBUFFER | W[0] = ALU, V[0].y = ALU.w   | EYE VECTOR                 |
| WLV0    | WBUFFER,VBUFFER | W[1] = ALU, V[1].y = ALU.w   | LIGHT0 DIRECTION VECTOR    |
| WLV1    | WBUFFER,VBUFFER | W[2] = ALU, V[2].y = ALU.w   | LIGHT1 DIRECTION VECTOR    |
| WLV2    | WBUFFER,VBUFFER | W[3] = ALU, V[3].y = ALU.w   | LIGHT2 DIRECTION VECTOR    |
| WLV3    | WBUFFER,VBUFFER | W[4] = ALU, V[4].y = ALU.w   | LIGHT3 DIRECTION VECTOR    |
| WLV4    | WBUFFER,VBUFFER | W[5] = ALU, V[5].y = ALU.w   | LIGHT4 DIRECTION VECTOR    |
| WLV5    | WBUFFER,VBUFFER | W[6] = ALU, V[6].y = ALU.w   | LIGHT5 DIRECTION VECTOR    |
| WLV6    | WBUFFER,VBUFFER | W[7] = ALU, V[7].y = ALU.w   | LIGHT6 DIRECTION VECTOR    |
| WLV7    | WBUFFER,VBUFFER | W[8] = ALU, V[8].y = ALU.w   | LIGHT7 DIRECTION VECTOR    |
| WSL0    | WBUFFER         | W[9] = ALU                   | SPOTLIGHT0 cos             |
| WSL1    | WBUFFER         | W[10] = ALU                  | SPOTLIGHT1 cos             |
| WSL2    | WBUFFER         | W[11] = ALU                  | SPOTLIGHT2 cos             |
| WSL3    | WBUFFER         | W[12] = ALU                  | SPOTLIGHT3 cos             |
| WSL4    | WBUFFER         | W[13] = ALU                  | SPOTLIGHT4 cos             |
| WSL5    | WBUFFER         | W[14] = ALU                  | SPOTLIGHT5 cos             |
| WSL6    | WBUFFER         | W[15] = ALU                  | SPOTLIGHT6 cos             |
| WSL7    | WBUFFER         | W[16] = ALU                  | SPOTLIGHT7 cos             |
| VED     | VBUFFER         | V[0].x = 1.0, V[0].z = ALU.w | EYE RADIAL DISTANCE VECTOR |
| VLD0    | VBUFFER         | V[1].x = 1.0, V[1].z = ALU.w | LIGHT0 DISTANCE VECTOR     |
| VLD1    | VBUFFER         | V[2].x = 1.0, V[2].z = ALU.w | LIGHT1 DISTANCE VECTOR     |
| VLD2    | VBUFFER         | V[3].x = 1.0, V[3].z = ALU.w | LIGHT2 DISTANCE VECTOR     |
| VLD3    | VBUFFER         | V[4].x = 1.0, V[4].z = ALU.w | LIGHT3 DISTANCE VECTOR     |
| VLD4    | VBUFFER         | V[5].x = 1.0, V[5].z = ALU.w | LIGHT4 DISTANCE VECTOR     |
| VLD5    | VBUFFER         | V[6].x = 1.0, V[6].z = ALU.w | LIGHT5 DISTANCE VECTOR     |
| VLD6    | VBUFFER         | V[7].x = 1.0, V[7].z = ALU.w | LIGHT6 DISTANCE VECTOR     |
| VLD7    | VBUFFER         | V[8].x = 1.0, V[8].z = ALU.w | LIGHT7 DISTANCE VECTOR     |
| VC0     | VBUFFER,TBUFFER | V[9] = ALU, T[1] = ALU       | DIFFUSE COLOR              |
| VC1     | VBUFFER,TBUFFER | V[10] = ALU, T[2] = ALU      | SPECULAR COLOR             |
| VNRM    | VBUFFER         | V[11] = ALU                  | NORMAL VECTOR              |
| VED2    | VBUFFER         | V[12] = ALU                  | EYE PLANAR DISTANCE VECTOR |
| TWV_NOP |                 |                              | NO VALID OUTPUT.           |

FIG. 10

| MICROCODE FIELD | BITS | LOCATION | DELAY | DESCRIPTION                 |
|-----------------|------|----------|-------|-----------------------------|
| oa              | 6    | 0:5      | 2     | OUTPUT BUFFER WRITE ADDRESS |
| rra             | 3    | 6:8      | 0     | RLU READ ADDRESS            |
| rwm             | 4    | 9:12     | 2     | RLU WRITE MASK              |
| rwa             | 3    | 13:15    | 2     | RLU WRITE ADDRESS           |
| ilu             | 2    | 16:17    | 2     | ILU OPERATION               |
| alu             | 4    | 18:21    | 1     | ALU OPERATION               |
| ais             | 2    | 22:23    | 1     | ALU SIGN CONTROL            |
| aia             | 1    | 24       | 1     | ALU INPUT A MUX             |
| mlu             | 3    | 25:27    | 0     | MLU OPERATION               |
| mib             | 2    | 28:29    | 0     | MLU INPUT B MUX             |
| mia             | 2    | 30:31    | 0     | MLU INPUT A MUX             |
| va              | 3    | 32:34    | 0     | INPUT BUFFER READ ADDRESS   |
| ce              | 1    | 35       | 0,2   | CONTEXT MEMORY READ/WRITE   |
| ca              | 6    | 36:41    | 0,2   | CONTEXT MEMORY ADDRESS      |
| mr              | 2    | 42:43    | 0     | MLU INPUT VECTOR ROTATE     |

FIG. 11




FIG. 12




FIG. 13




FIG. 14






FIG. 14A




FIG. 14B

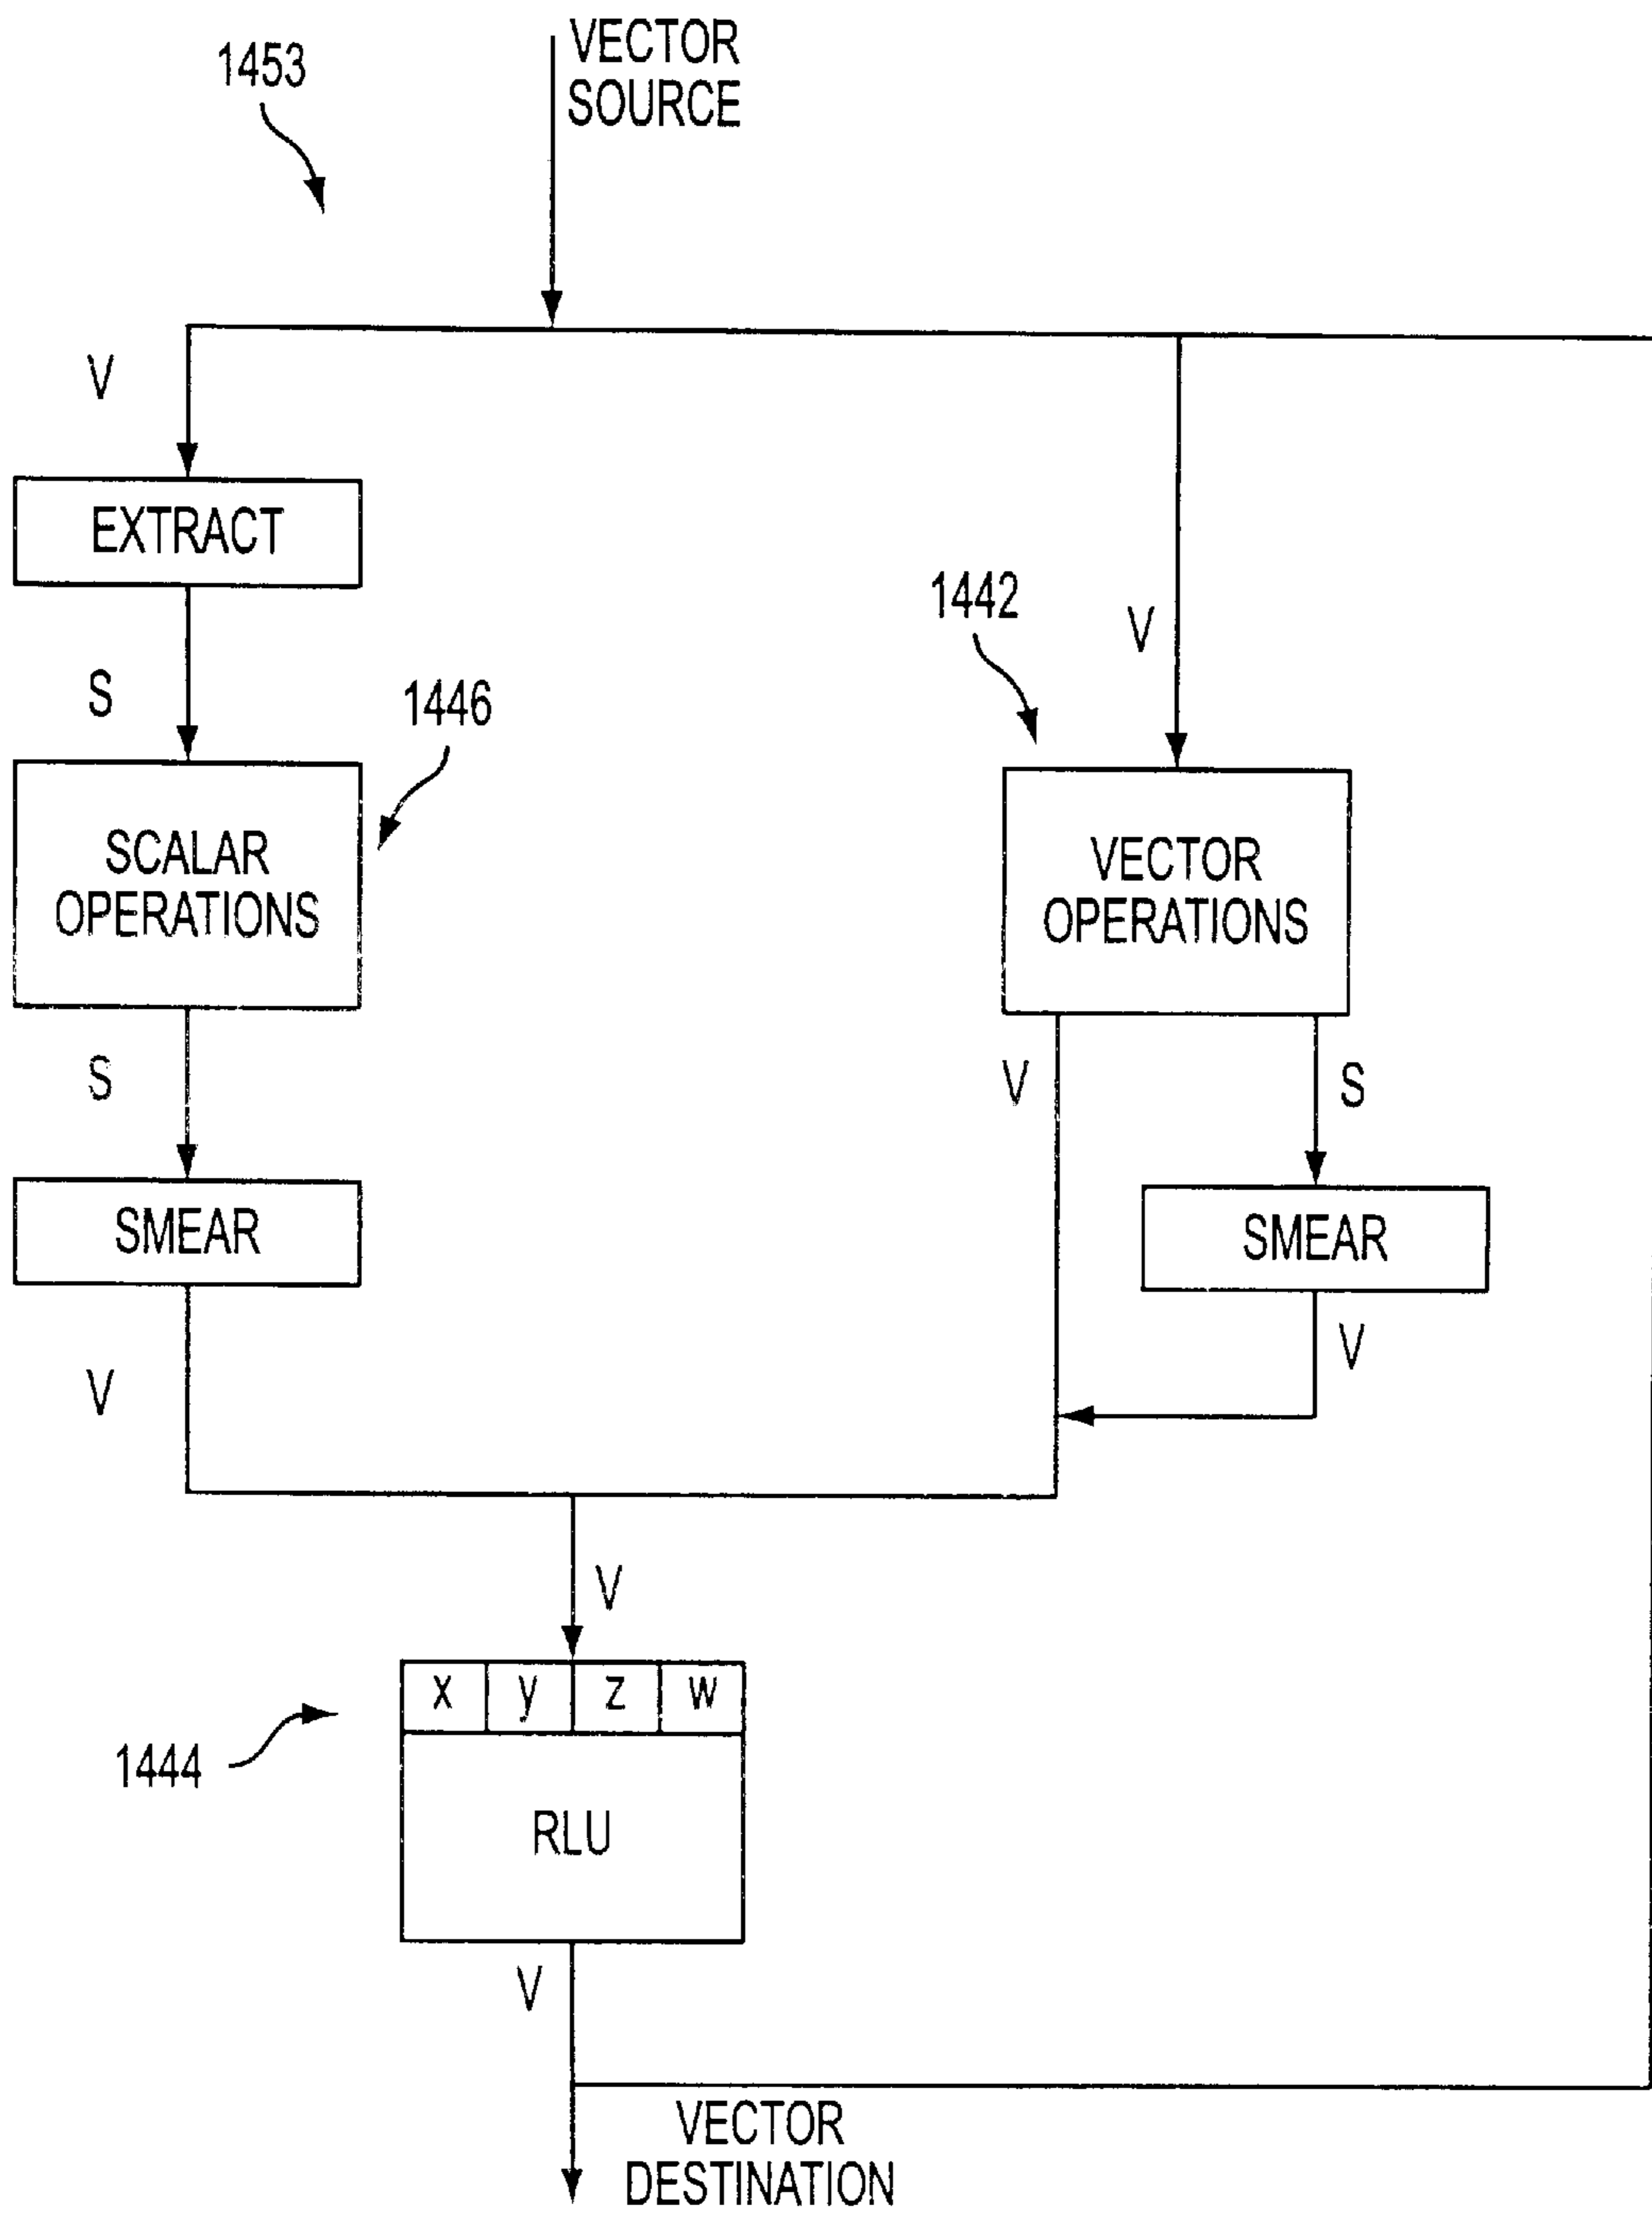


FIG. 14C




FIG. 14D




FIG. 15






FIG. 16




FIG. 17




FIG. 18




FIG. 19




FIG. 20

| NAME | REGISTER | DESCRIPTION  |
|------|----------|--|
| Z    | IFLAG    | CLEAR FLAG. SETS IFLAG AND MFLAG TO 0.   |
| C    | IFLAG    | SPOTLIGHT CONE FLAG. SET IF VERTEX IS OUTSIDE SPOTLIGHT CONE.                                |
| S    | IFLAG    | SPECULAR2 FLAG. SET IF SPECULAR CONTRIBUTION IS NEGATIVE.                                    |
| D    | IFLAG    | DIFFUSE FLAG. SET IF DIFFUSE TERM IS NEGATIVE.   |
|      | MFLAG    |  |
| U    | MFLAG    | SPOTLIGHT CONE ATTENUATION FLAG. SET IF SPOTLIGHT CONE ATTENUATION CONTRIBUTION IS NEGATIVE. |
| T    | MFLAG    | SPECULAR FLAG. SET IF SPECULAR CONTRIBUTION IS NEGATIVE.                                     |
| R    | MFLAG    | RANGE FLAG. SET IF VERTEX IS TOO FAR AWAY FROM THE LIGHT.                                    |

FIG. 21

| MICROCODE FIELD | BITS | LOCATION | DELAY | DESCRIPTION                      |
|-----------------|------|----------|-------|----------------------------------|
| oa              | 3    | 0:2      | 2     | OUTPUT ADDRESS                   |
| rwe             | 1    | 3        | 2     | RLU WRITE ENABLE                 |
| rwa             | 2    | 4:5      | 2     | RLU WRITE ADDRESS                |
| R23             | 1    | 6        | 0     | RLU (MLU) READ ADDRESS           |
| R01             | 1    | 7        | 1     | RLU (ALU) READ ADDRESS           |
| aia             | 1    | 8        | 1     | ALU INPUT A MUX                  |
| alu             | 2    | 9:10     | 1     | ALU OPERATION                    |
| mib             | 2    | 11:12    | 0     | MLU INPUT B MUX                  |
| mia             | 2    | 13:14    | 0     | MLU INPUT A MUX                  |
| mlu             | 2    | 15:16    | 0     | MLU OPERATION                    |
| mwa             | 5    | 17:21    | 0     | MLU WBUFFER READ ADDRESS         |
| awa             | 5    | 22:26    | 1     | ALU WBUFFER READ ADDRESS         |
| va              | 4    | 27:30    | 0     | VBUFFER READ ADDRESS             |
| os              | 2    | 31:32    | 2     | LLU OUTPUT ADDRESS               |
| frm             | 6    | 33:38    | 2     | FLAG REGISTER MASK               |
| mfe             | 1    | 39       | 2     | MFLAG WRITE ENABLE               |
| mfa             | 2    | 40:41    | 2     | MFLAG WRITE ADDRESS              |
| ife             | 1    | 42       | 2     | IFLAG WRITE ENABLE               |
| ifa             | 2    | 43:44    | 2     | IFLAG WRITE ADDRESS              |
| fia             | 2    | 45:46    | 2     | FLU INPUT A MUX                  |
| flu             | 3    | 47:49    | 2     | FLU OPERATION                    |
| M1c             | 1    | 50       | 2     | MAC1 INPUT C MUX                 |
| M1b             | 2    | 51:52    | 2     | MAC1 INPUT B MUX                 |
| M1a             | 2    | 53:54    | 2     | MAC1 INPUT A MUX                 |
| M0c             | 2    | 55:56    | 2     | MAC0 INPUT C MUX                 |
| M0b             | 2    | 57:58    | 2     | MAC0 INPUT B MUX                 |
| M0a             | 2    | 59:60    | 2     | MAC0 INPUT A MUX                 |
| ce              | 3    | 61:63    | 0,2   | CONTEXT MEMORY READ/WRITE ENABLE |
| ca              | 6    | 64:69    | 0,2   | CONTEXT MEMORY ADDRESS           |
| C3a             | 4    | 70:73    | 2     | CONTEXT3 MEMORY ADDRESS          |
| C2a             | 4    | 74:77    | 2     | CONTEXT2 MEMORY ADDRESS          |
| C1a             | 5    | 78:82    | 2     | CONTEXT1 MEMORY ADDRESS          |
| C0a             | 2    | 83:84    | 2     | CONTEXT0 MEMORY ADDRESS          |

FIG. 22






FIG. 23




FIG. 24




FIG. 25




FIG. 25B

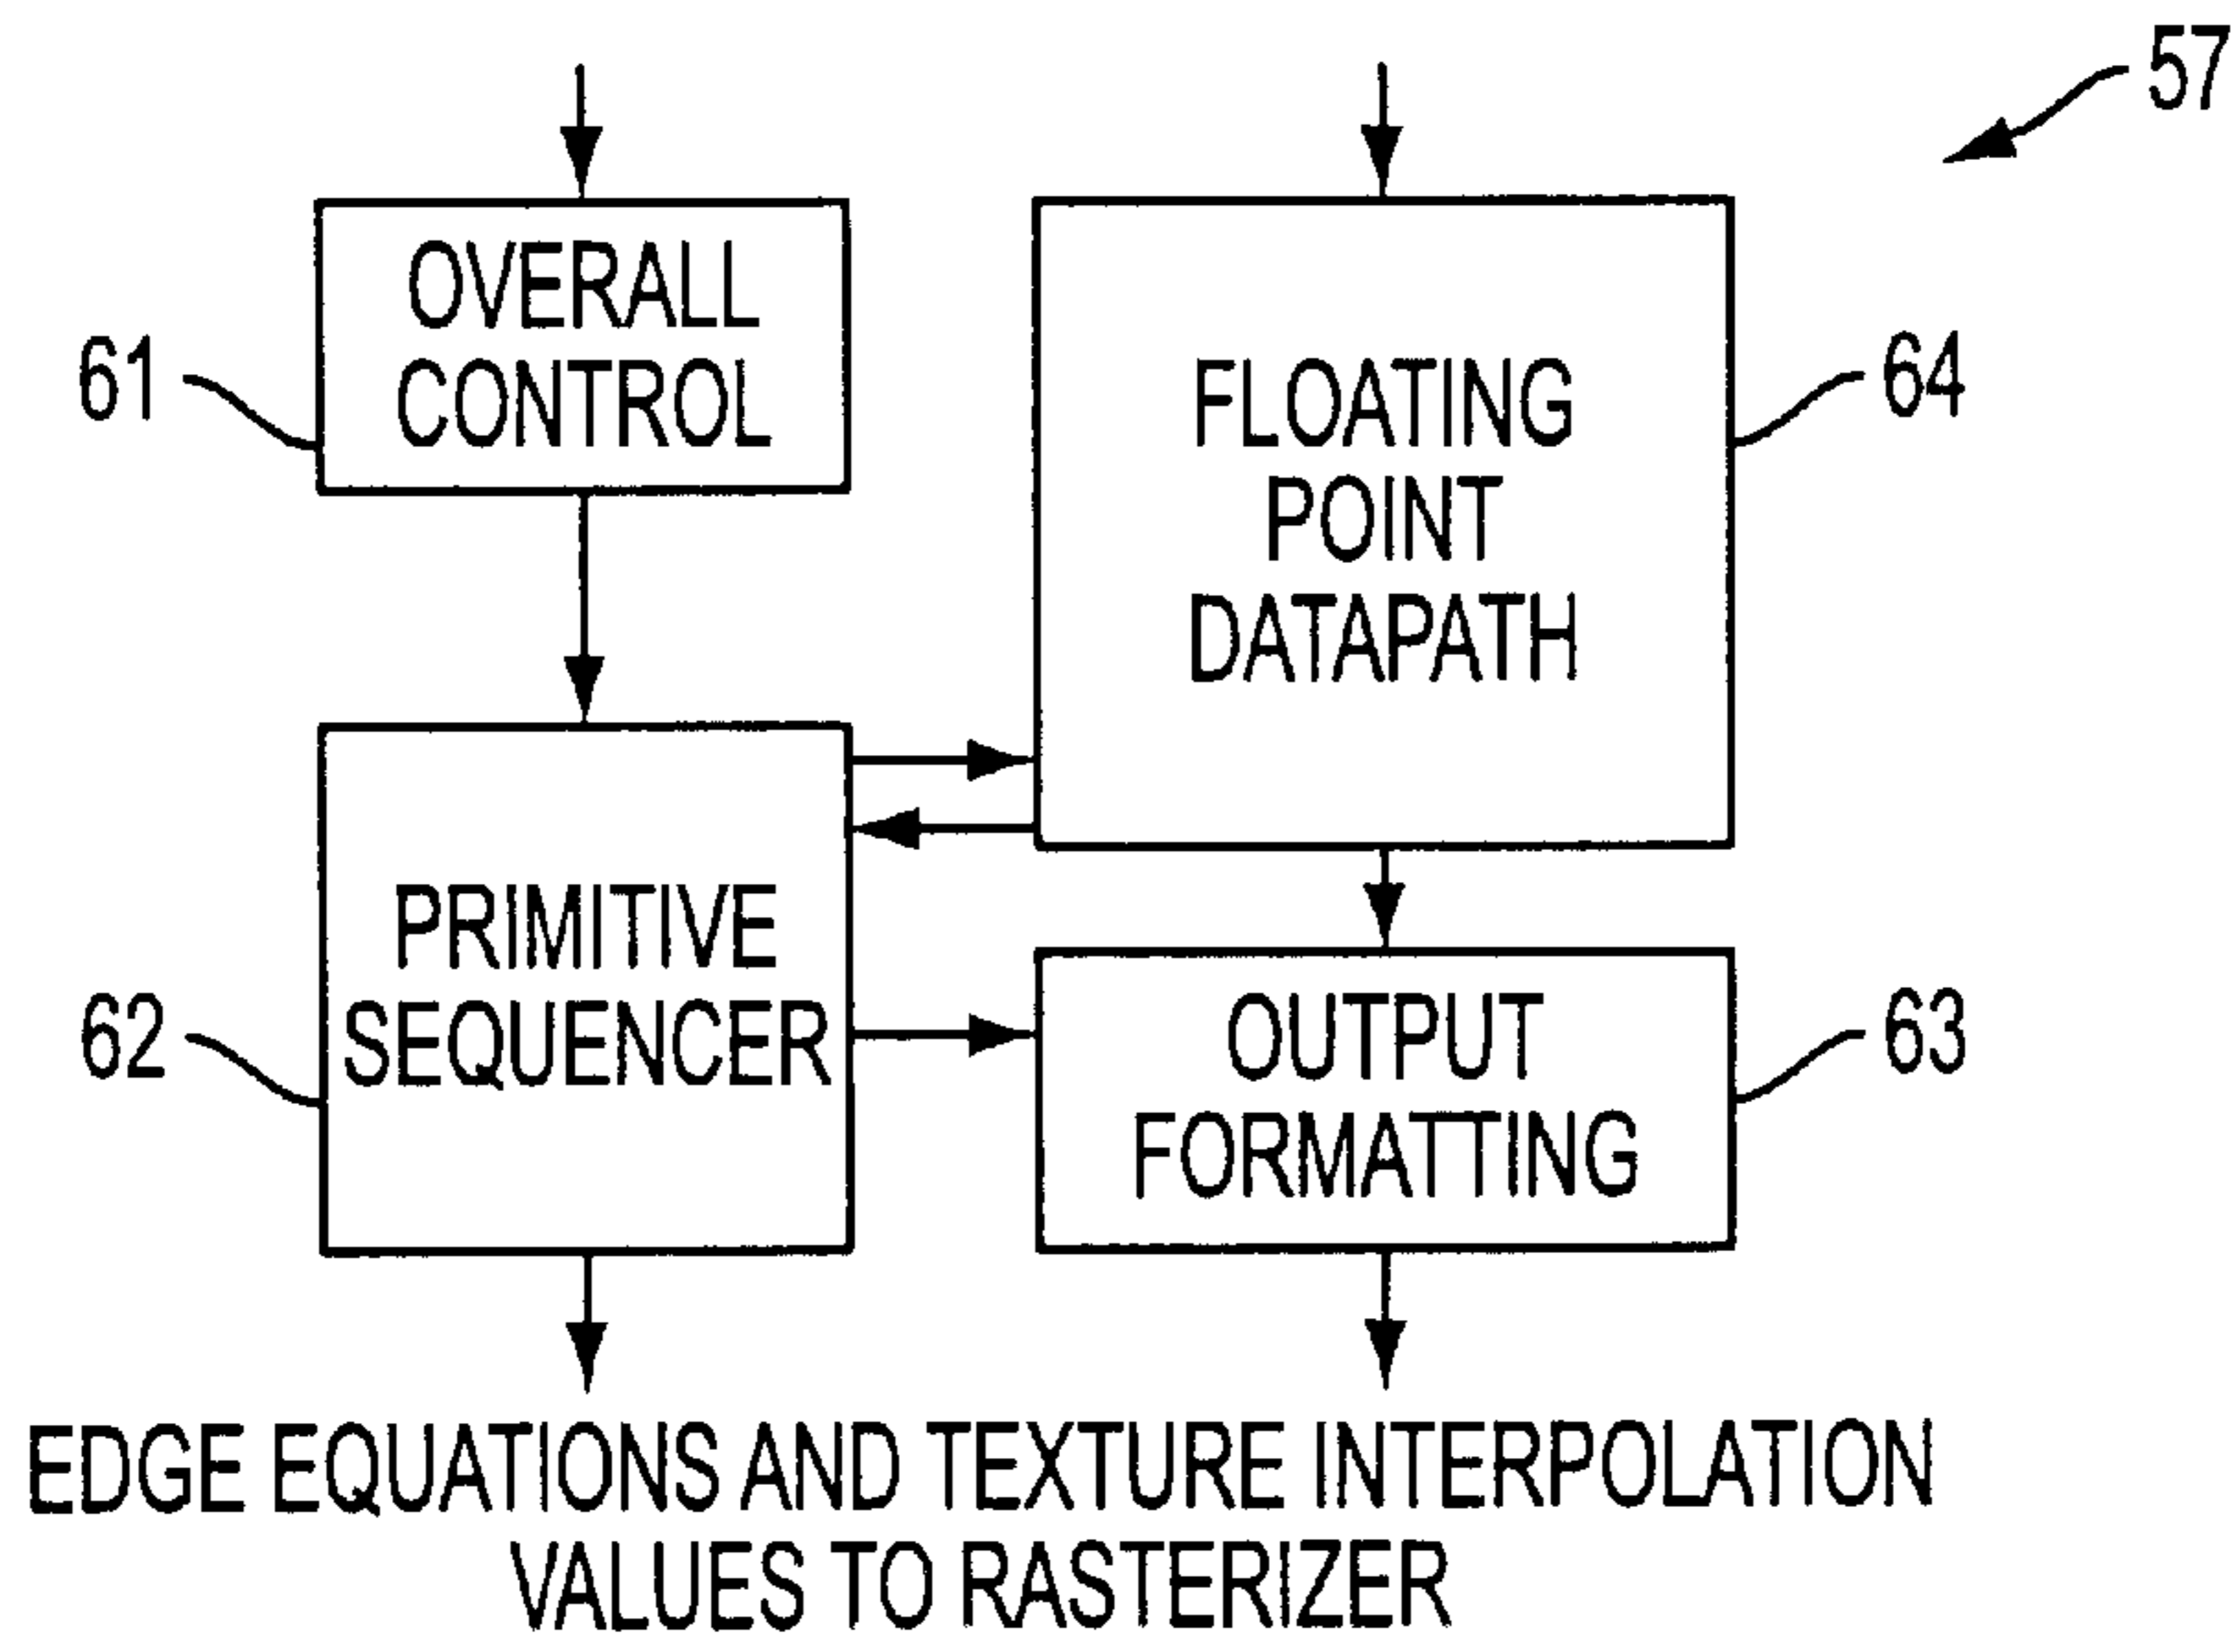


FIG. 26




FIG. 26A

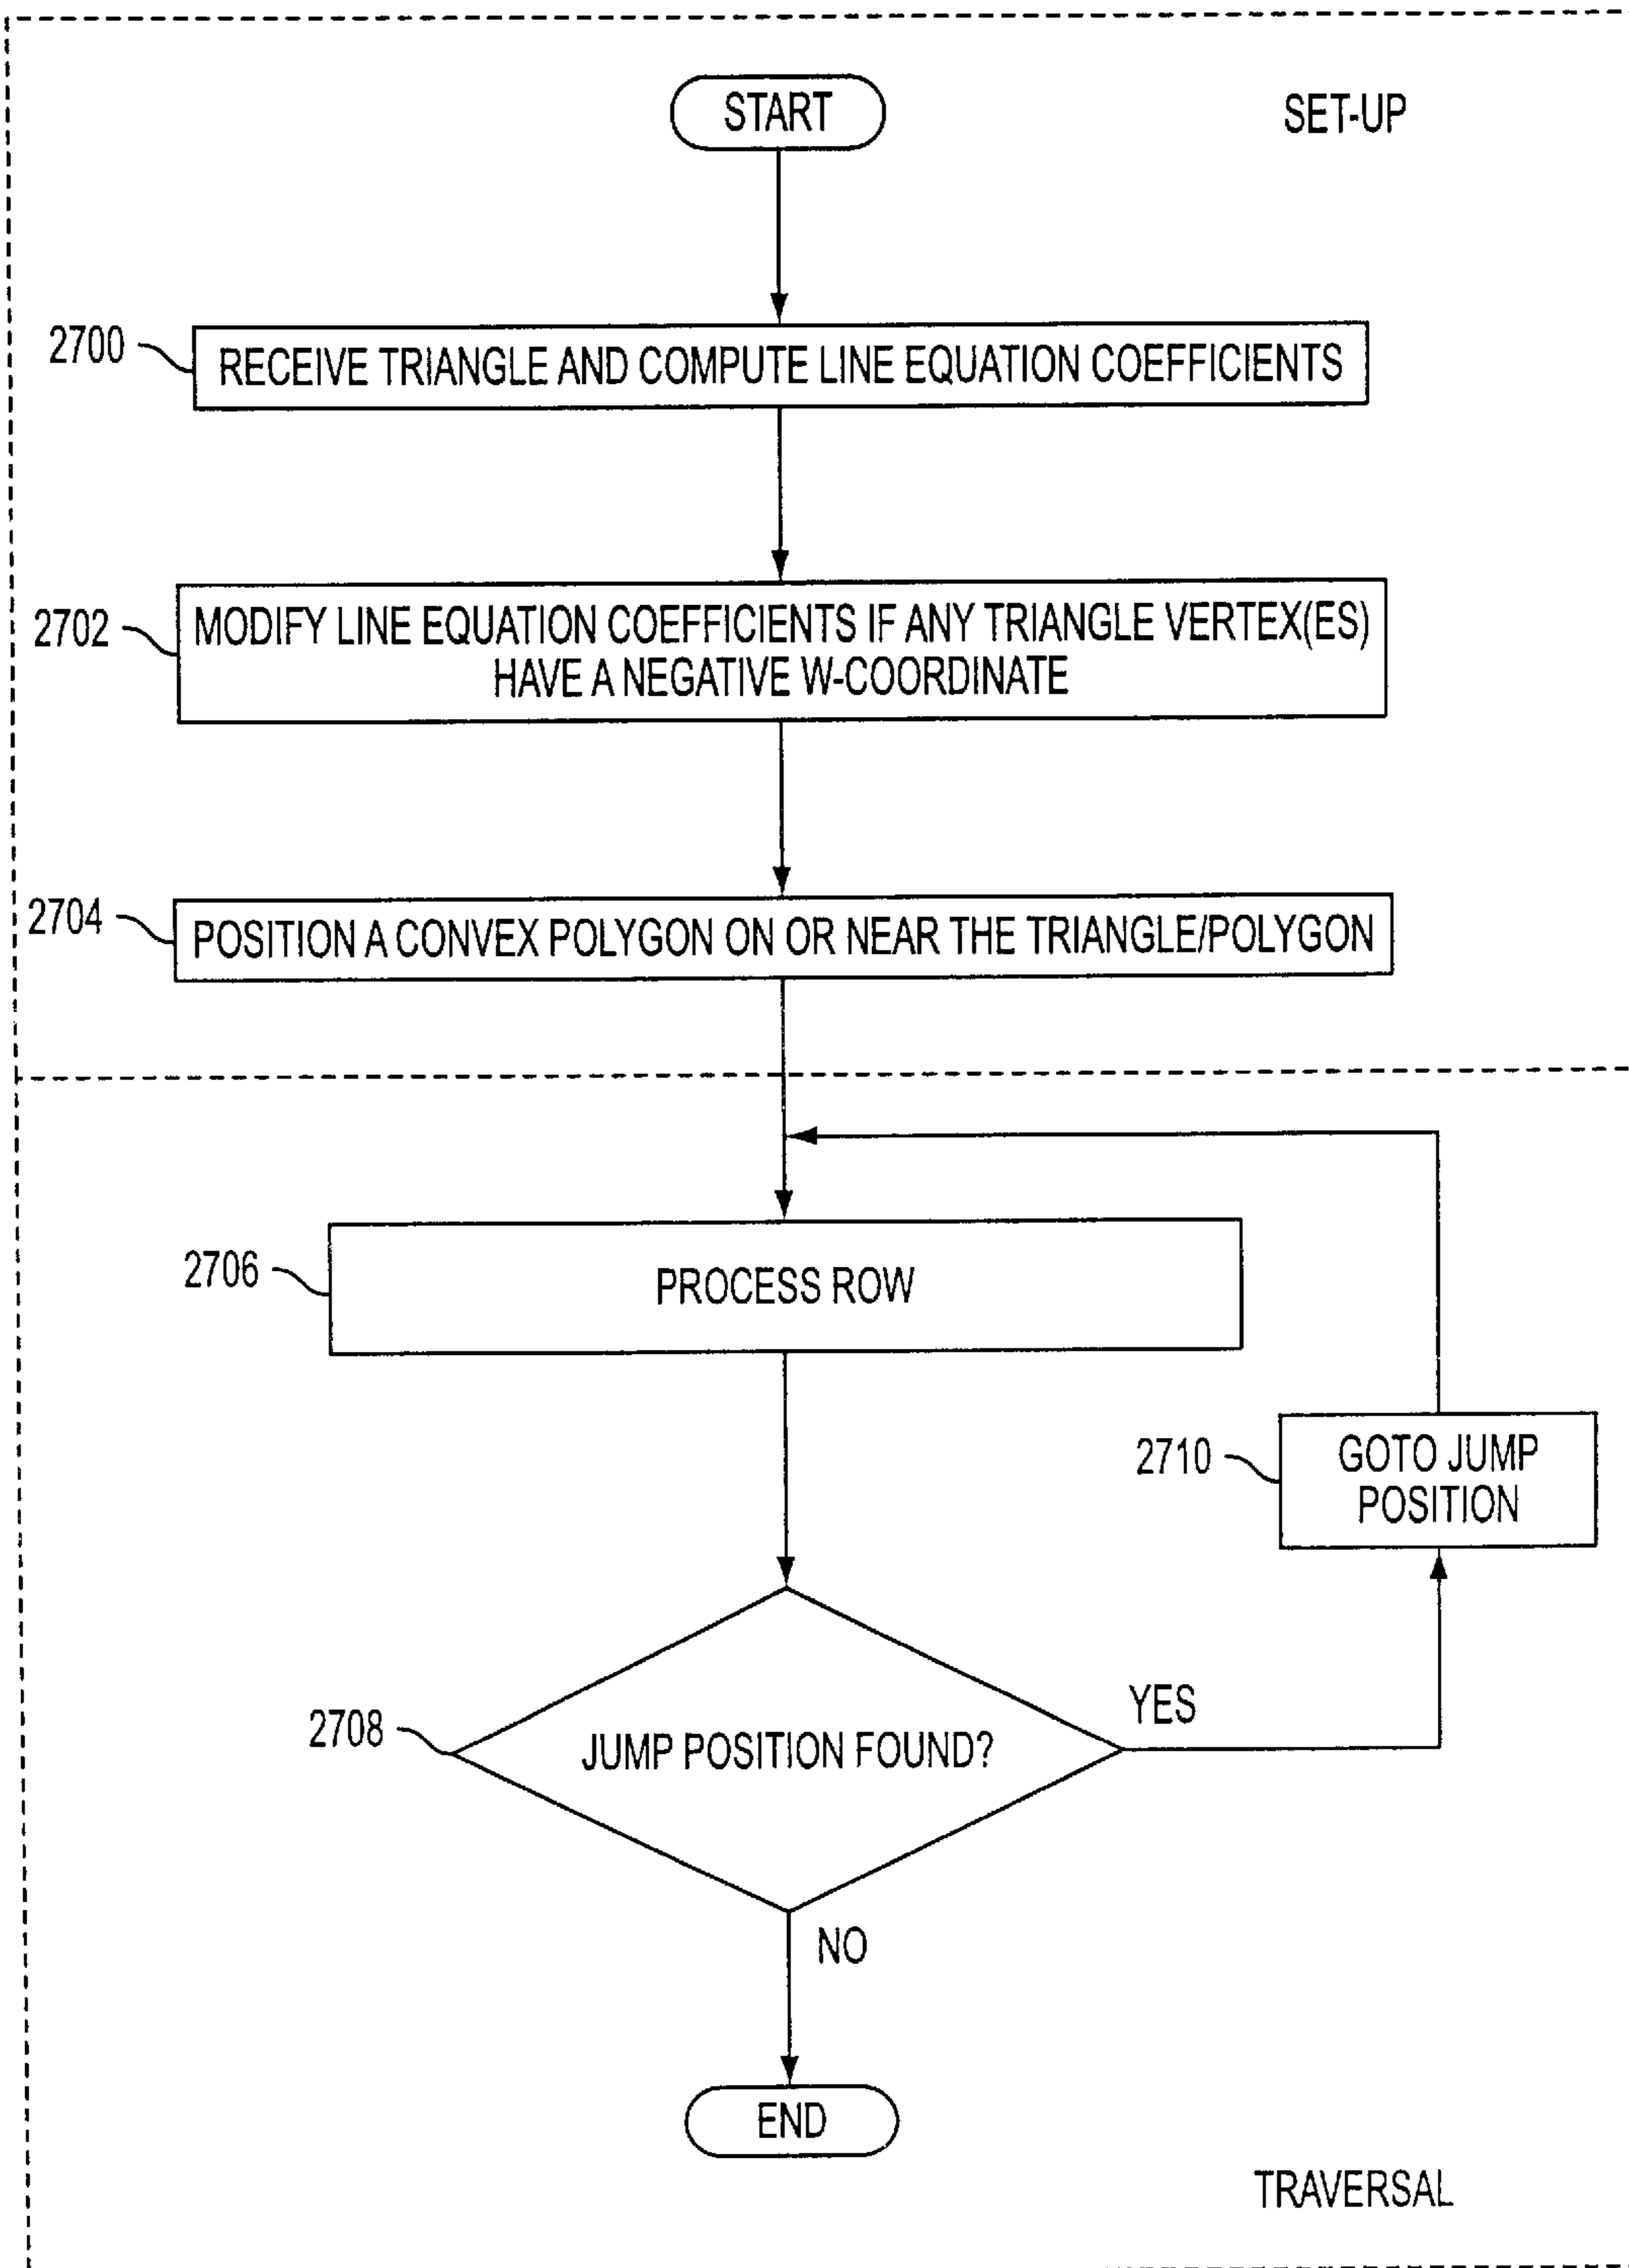


FIG. 27




FIG. 27A






FIG. 28




FIG. 28A




FIG. 28B

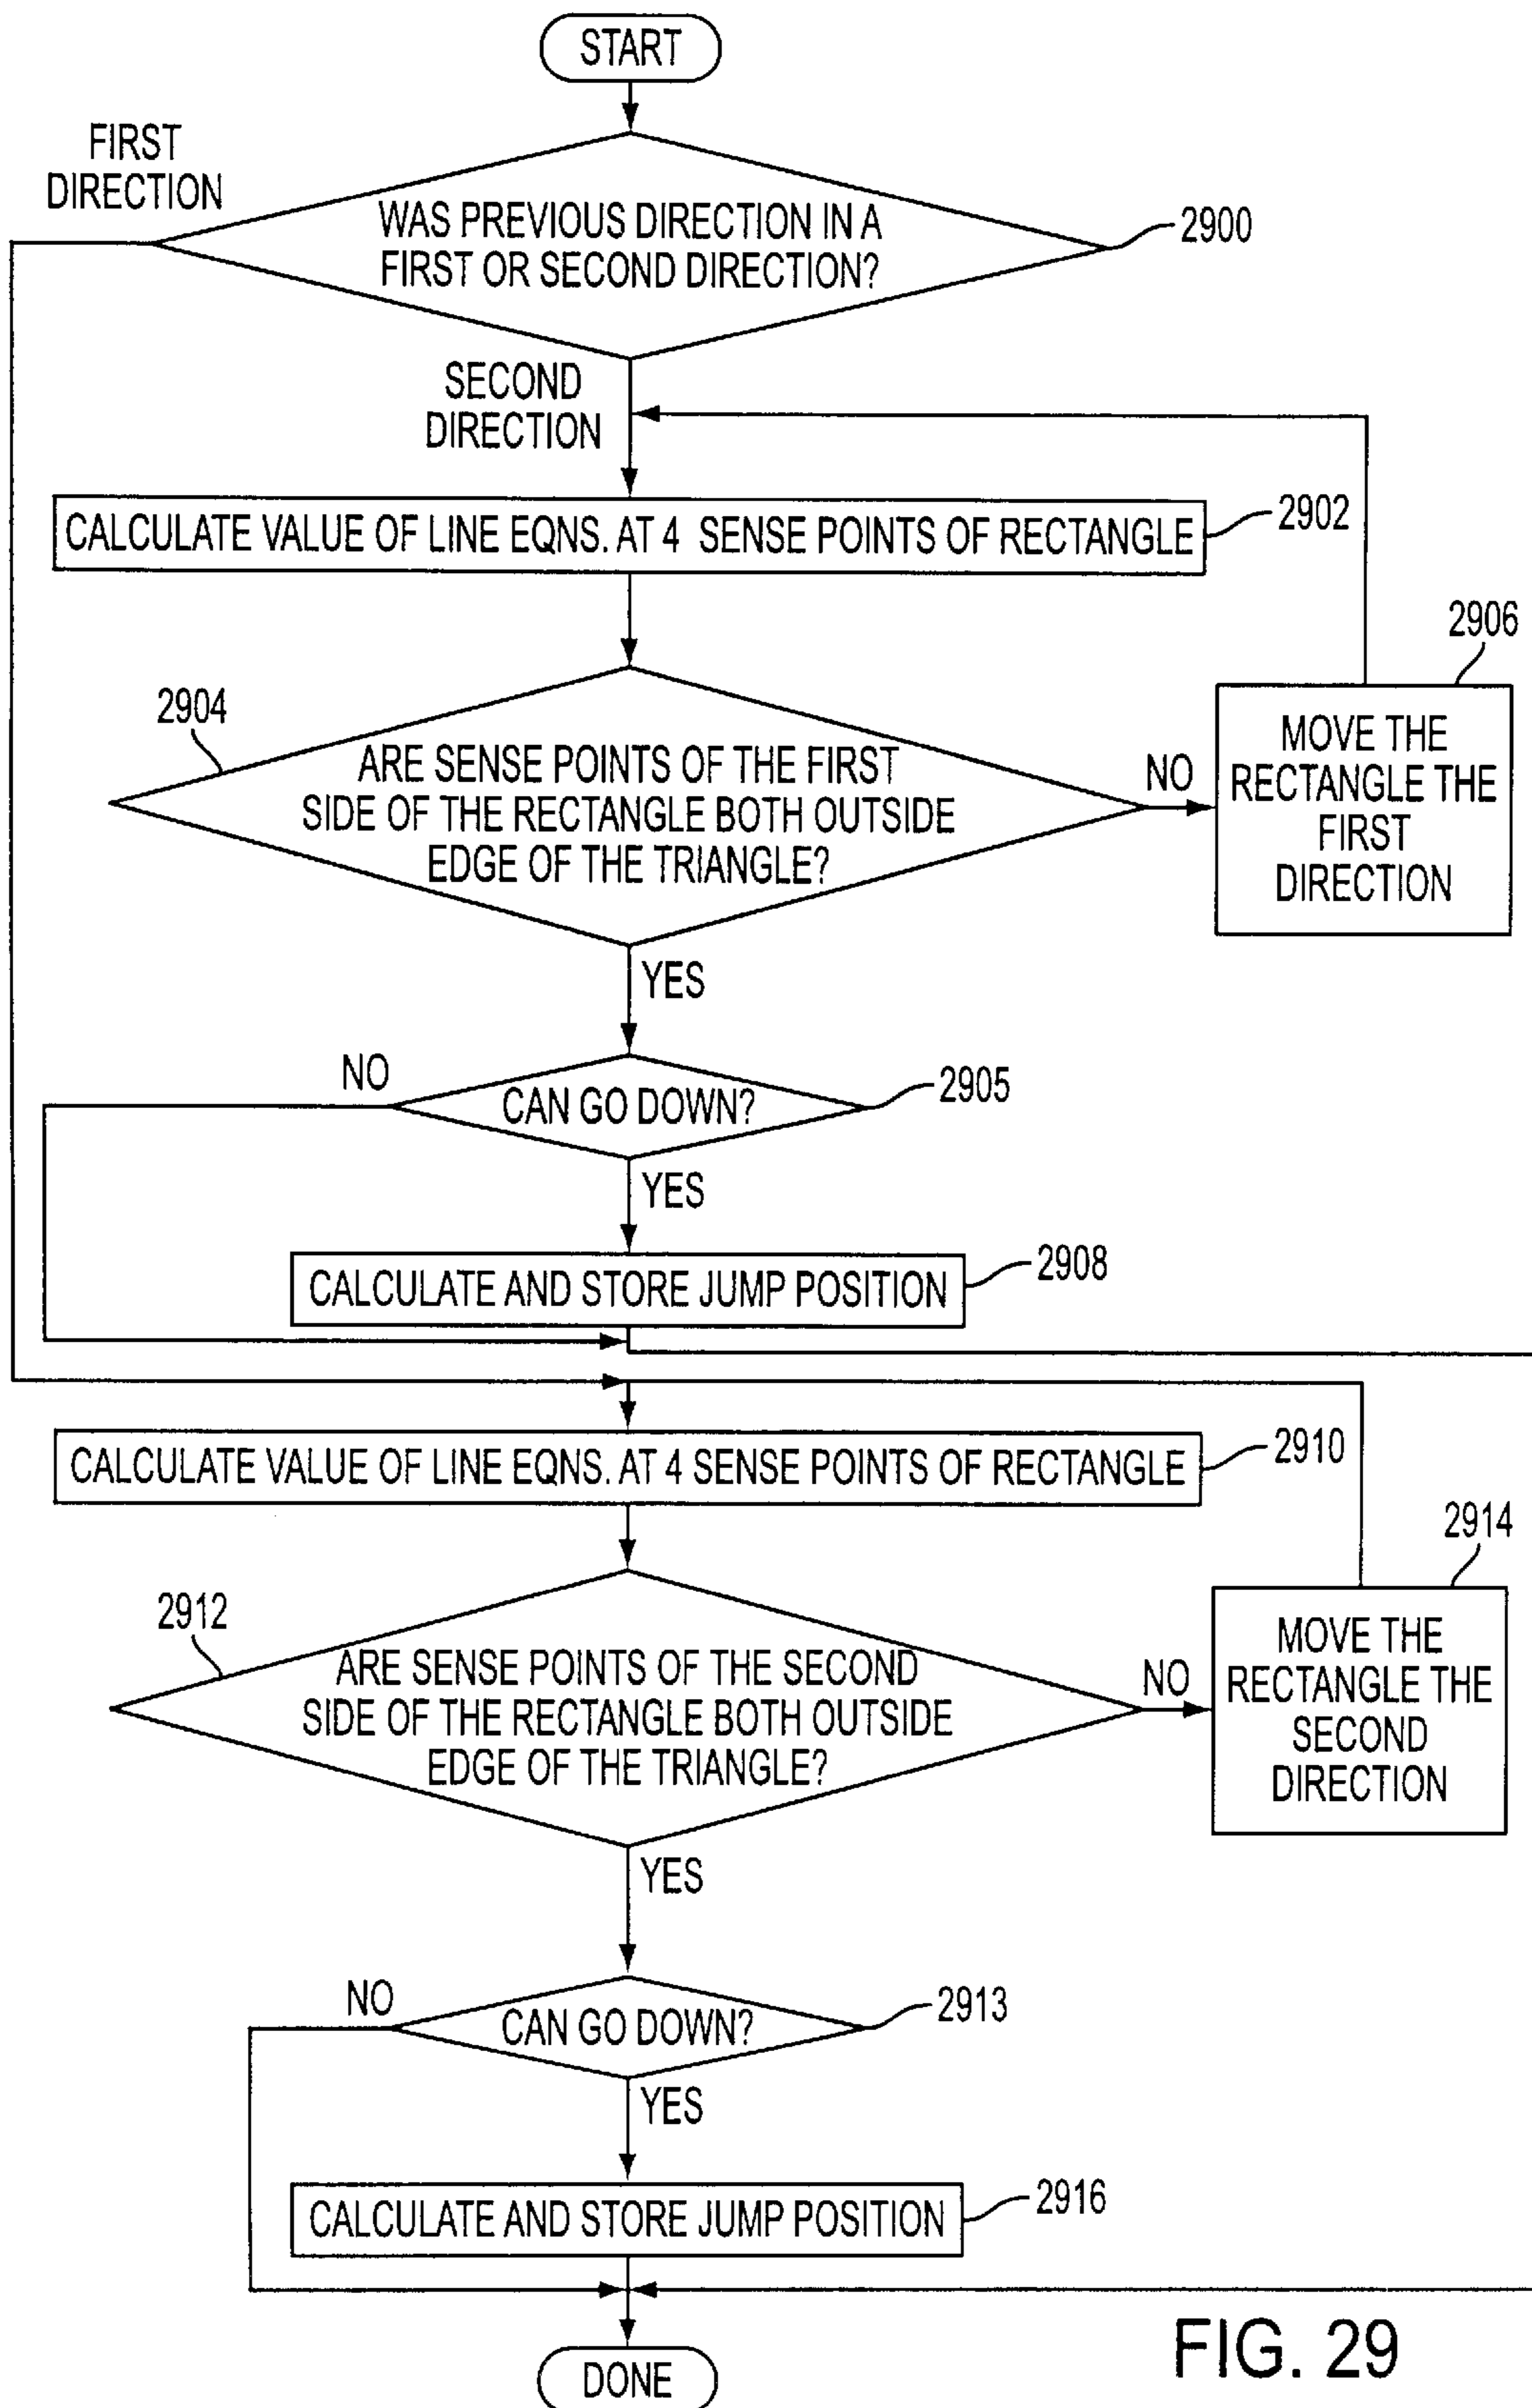


FIG. 29




FIG. 29A




FIG. 30

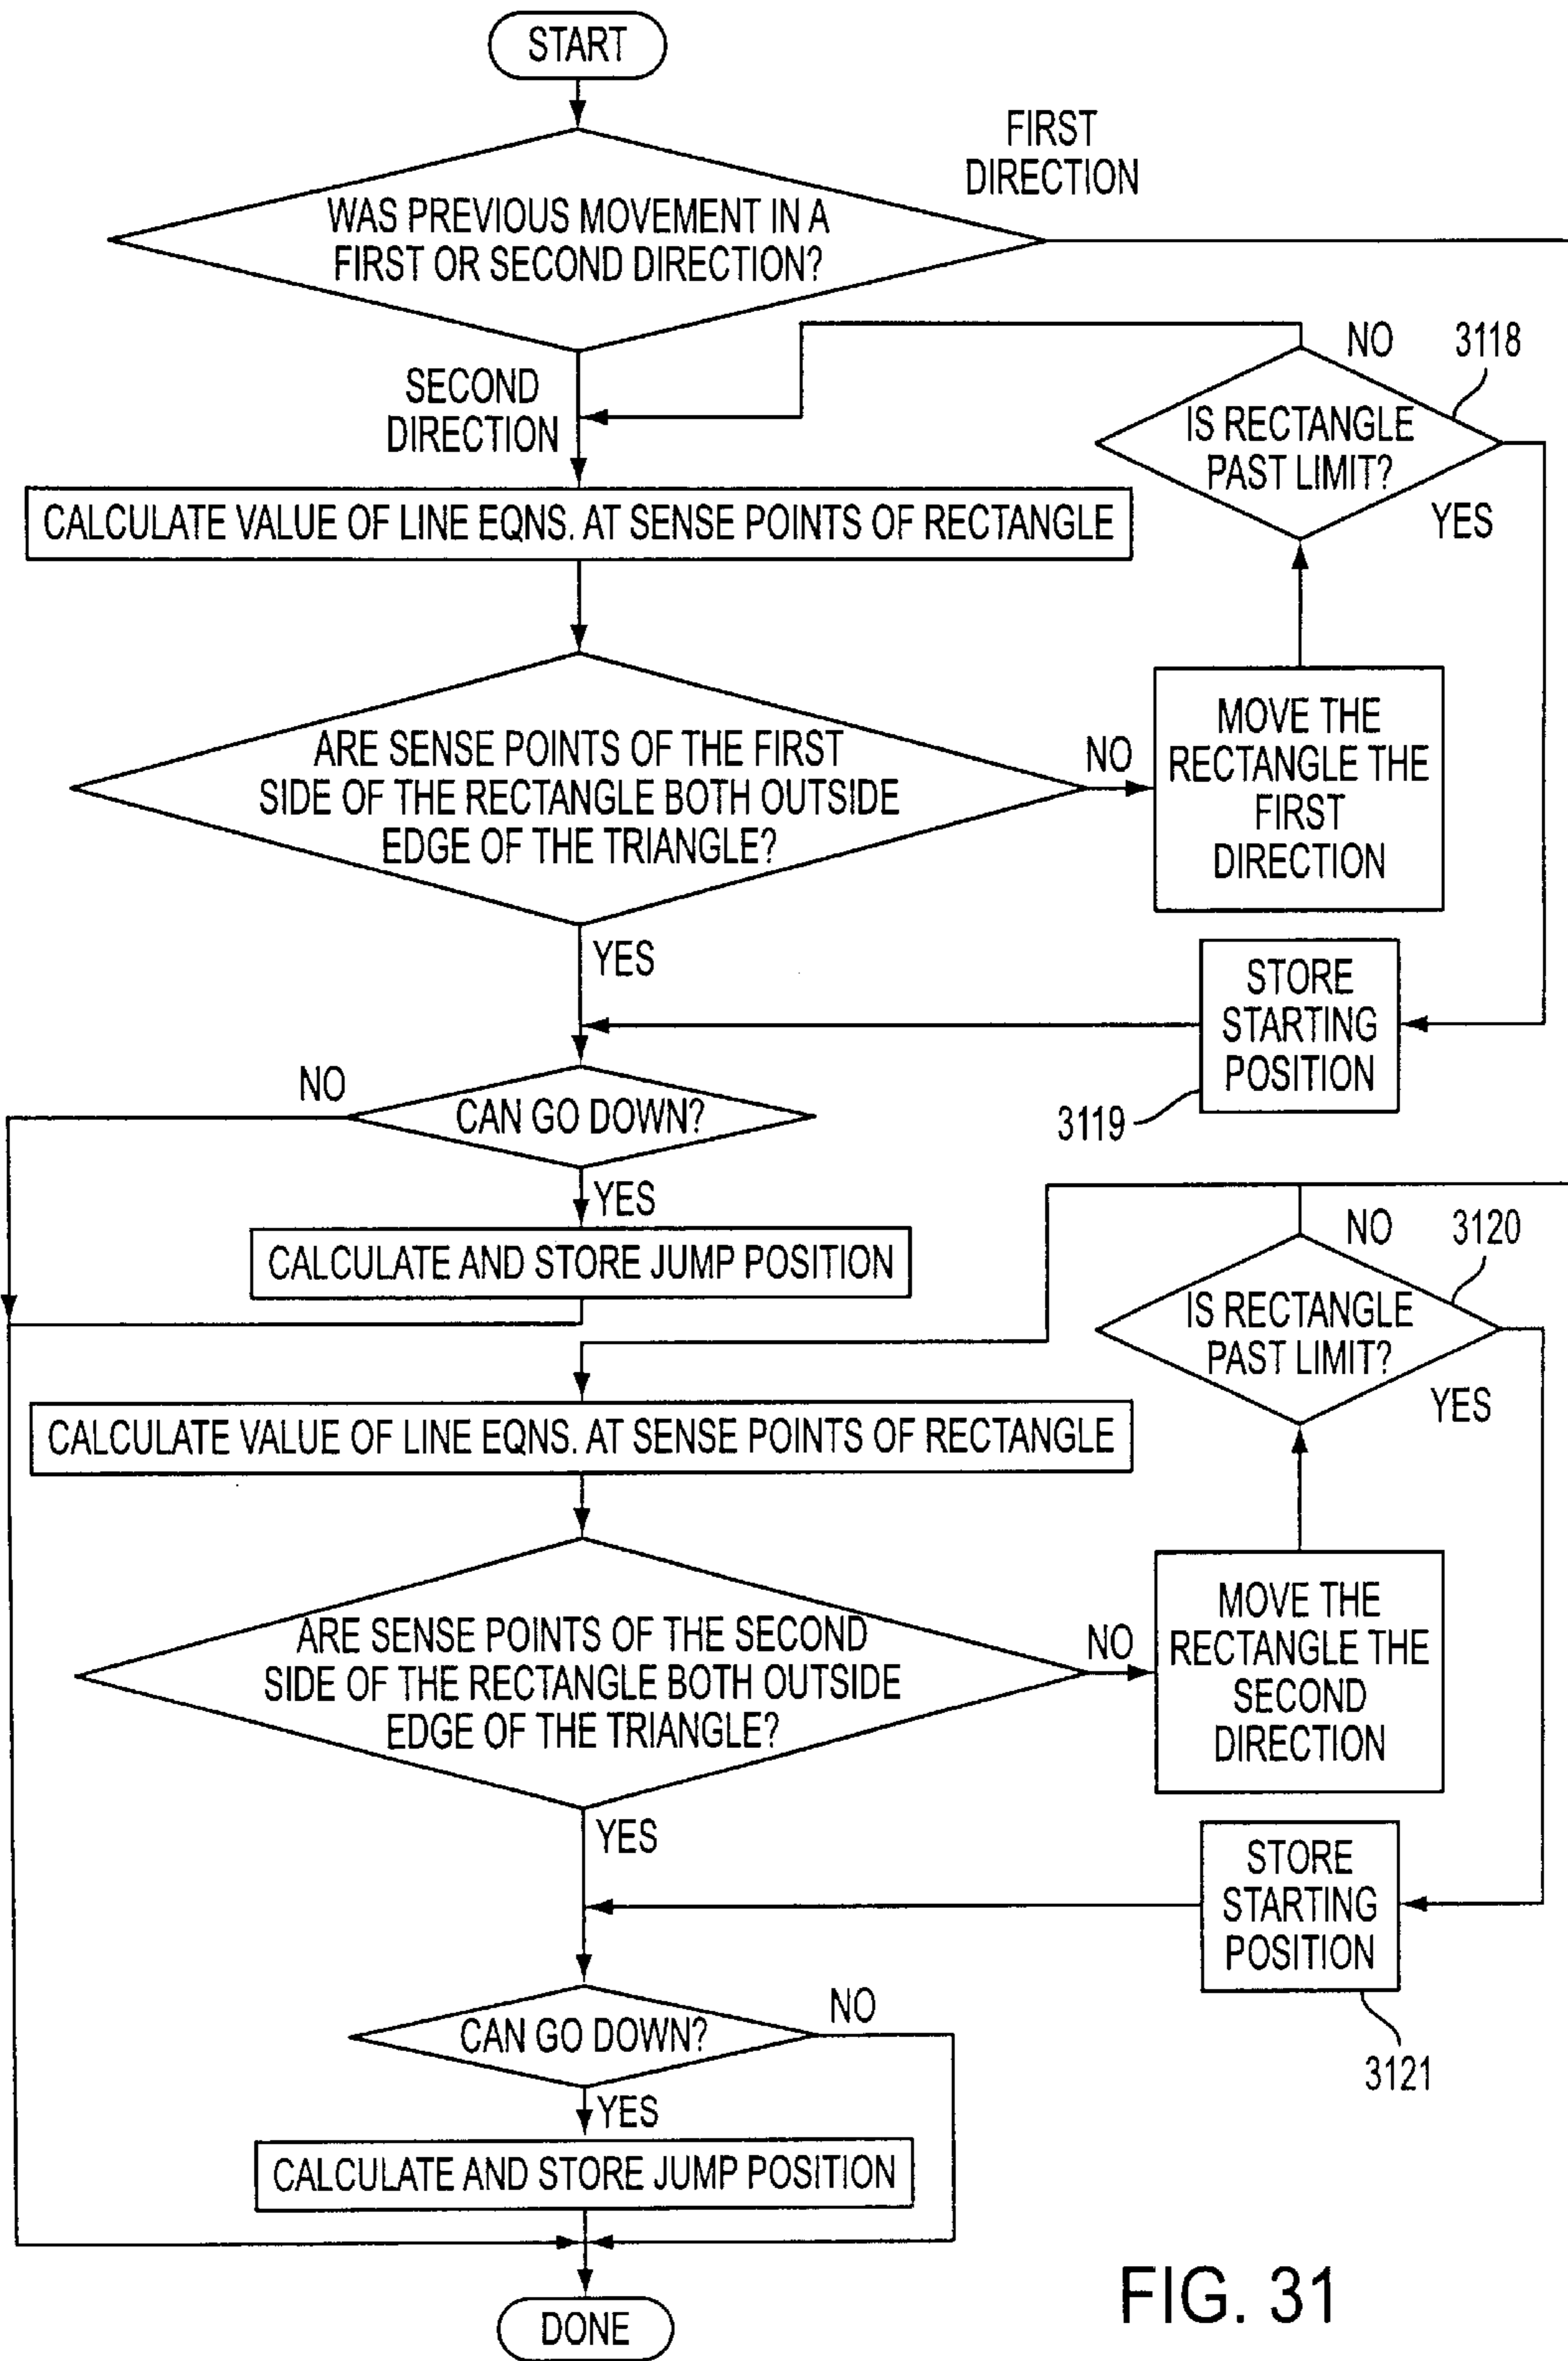


FIG. 31










FIG. 32




FIG. 32A




FIG. 32B




FIG. 32C

## METHOD, APPARATUS AND ARTICLE OF MANUFACTURE FOR BOUSTROPHEDONIC RASTERIZATION

### RELATED APPLICATIONS

The present application is related to applications entitled "Method, Apparatus and Article of Manufacture for Area Rasterization using Sense Points" which was filed on Dec. 6, 1999 under Ser. No. 09/455,305, "Method, Apparatus and Article of Manufacture for Clip-less Rasterization using Line Equation-based Traversal" which was filed on Dec. 6, 1999 under Ser. No. 09/455,728, "Transform, Lighting and Rasterization System Embodied on a Single Semiconductor Platform" which was filed on Dec. 6, 1999 under Ser. No. 09/454,516, and issued under U.S. Pat. No. 6,198,488, "Method, Apparatus and Article of Manufacture for a Vertex Attribute Buffer in a Graphics Processor" which was filed on Dec. 6, 1999 under Ser. No. 09/454,525, and issued under U.S. Pat. No. 6,515,671, "System, Method and Computer Program Product for a Blending Operation in a Transform Module of a Computer Graphics Pipeline" which was filed on Dec. 6, 1999 under Ser. No. 09/456,102, and issued under U.S. Pat. No. 6,353,439, "Method and Apparatus for a Lighting Module in a Graphics Processor" which was filed on Dec. 6, 1999 under Ser. No. 09/454,524, and issued under U.S. Pat. No. 6,417,851, and "Method, Apparatus and Article of Manufacture for a Sequencer in a Transform/Lighting Module Capable of Processing Multiple Independent Execution Threads" which was filed on Dec. 6, 1999 under Ser. No. 09/456,104, which were filed concurrently herewith, and which are all incorporated herein by reference in their entirety.

### FIELD OF THE INVENTION

The present invention relates generally to rasterizers and, more particularly, to optimizing results during the conversion of primitives defined by vertexes to equivalent images composed of pixel patterns that can be stored and manipulated as sets of bits.

### BACKGROUND OF THE INVENTION

Raster displays are commonly used in computer graphics systems. These displays store graphics images as a matrix of picture elements or "pixels" with data representing each pixel being stored in a display buffer. This data specifies the display attributes for each pixel on the screen such as the intensity and color of the pixel. An entire image is read from the display buffer and displayed on the screen by sequentially scanning out horizontal rows of pixel data or "scan lines."

Raster display systems commonly use polygons as basic building blocks or "primitives" for drawing more complex images. Triangles are a common basic primitive for polygon drawing systems, since a triangle is the simplest polygon and more complex polygons can always be represented as sets of triangles. The process of drawing triangles and other geometric primitives on the screen is known as "rasterization."

An important part of rasterization involves determining which pixels fall within a given triangle. Rasterization systems generally step from pixel to pixel in various ways and determine whether or not to "render," i.e. draw into a frame buffer or pixel map, each pixel as part of the triangle. This, in turn, determines how to set the data in the display buffer representing each pixel. Various traversal algorithms have been developed for moving from pixel to pixel in a way such that all pixels within the triangle are covered.

Rasterization systems sometimes represent a triangle as a set of three edge-functions. An edge function is a line equation representing a straight line, which serves to subdivide a two-dimensional plane. Edge functions classify each point within the plane as falling into one of three regions: the region "inside" of the triangle, the region "outside" of the triangle, or the region representing a line itself. The type of edge function that will be discussed has the property that points "inside" of the triangle have a value greater than zero, points "outside" have a value less than zero, and points exactly on the line have a value of zero. This is shown in Prior Art FIG. 1. Applied to rasterization systems, the two-dimensional plane is represented by the graphics screen, points are represented by individual pixels, and the edge function serves to subdivide the graphics screen.

Triangles are created by the union of three edges, or more particularly three half-planes, each of which is specified by edge functions. It is possible to define more complex polygons by using Boolean combinations of more than three edges. Since the rasterization of triangles involves determining which pixels to render, a tie-breaker rule is generally applied to pixels that lie exactly on any of the edges to determine whether the pixels are to be considered interior or exterior to the triangle.

As shown in Prior Art FIG. 1A, each pixel has associated with it a set of edge variables, ( $e_0$ ,  $e_1$  and  $e_2$ ), which are proportional to the signed distance between the pixel and the three respective edges. The value of each edge variable is determined for a given triangle by evaluating the three edge functions,  $f_0(x,y)$ ,  $f_1(x,y)$  and  $f_2(x,y)$  for the pixel location. It is important to note that it can be determined whether or not a pixel falls within a triangle by looking at only the signs of  $e_0$ ,  $e_1$  and  $e_2$ .

In determining which pixels to render within a triangle, typical rasterization systems compute the values of the edge variables, ( $e_0$ ,  $e_1$  and  $e_2$ ), for a given set of three edge functions and a given pixel position, and then use a set of increment values ( $\Delta e_{outside}$ ,  $\Delta e_{inside}$ , etc.) to determine the edge variable values for adjacent pixels. The rasterization system traverses the triangle, adding the increment values to the current values as a traversal algorithm steps from pixel to pixel.

With reference again to Prior Art FIG. 1, a line is illustrated that is defined by two points: (X,Y) and (X+dX, Y+dY). As noted above, this line can be used to divide the two dimensional space into three regions: all points "outside" of, "inside" of, and exactly on the line.

The edge  $f(x,y)$  can be defined as:

$$f(x,y)=(x-X)dY-(y-Y)dX.$$

This function has the useful property that its value is related to the position of the point (x,y) relative to the edge defined by the points (X,Y) and (X+dX, Y+dY):

$f(x,y)>0$  if (x,y) is "inside";

$f(x,y)=0$  if (x,y) is exactly on the line; and

$f(x,y)<0$  if (x,y) is "outside".

Existing rasterization systems commonly use this function, since it can be computed incrementally by simple addition:

$$f(x+1,y)=f(x,y)+dY; \text{ and}$$

$$f(x,y+1)=f(x,y)-dX.$$

A variety of different traversal algorithms are presently used by different rasterization systems in the rendering



process. Any algorithm guaranteed to cover all of the pixels within the triangle can be used.

For example, some solutions involve generating the pixels in a unidirectional manner. Such traditional unidirectional solution includes generating the pixels row-by-row in a constant direction. This requires that the sequence shift across the primitive to a starting location on a first side of the primitive upon finishing at a location on an opposite side of the primitive. Each time this shift is executed, pixels or texture values are stored which were not positioned adjacent to pixels or texture values processed immediately beforehand. Therefore, such distant pixels or texture values have a greater chance of belonging to different memory access blocks, making such access inefficient.

This problem is exacerbated with primitives of a large size. For example, if a sequence generates pixels along a very long row or column, the probability is reduced that contiguous pixels or texture values from an adjacent row or column, respectively, will still be in memory. This probability is further reduced with memory of a limited size.

As yet another example, a haphazard generation sequence includes jumping around randomly within the primitive. Using this method, pixels or texture values that are accessed are rarely adjacent to those that were accessed recently enough to still be in a limited size memory.

There is therefore a need for a rasterization system that processes pixels or texture values that are frequently adjacent to pixels or texture values that have been previously processed immediately beforehand and are still stored in memory.

#### DISCLOSURE OF THE INVENTION

A method, apparatus and article of manufacture are provided for performing rasterization using alternating sense point traversal. Upon receipt of a primitive, i.e. a triangle, a plurality of points is positioned on or near the primitive. Such points define an enclosed convex region and are located at corners of the convex region. In operation, the points and thus the convex region are moved in an alternating manner for the purpose of identifying an area in the primitive for rendering pixels therein. In particular, the points are moved in a boustrophedonic manner.

This boustrophedonic rasterization constrains the sequence to obey certain rules that offer better performance for hardware. Boustrophedonic refers to a serpentine pattern that folds back and forth. A horizontal boustrophedonic sequence, for example, may generate all the pixels within a primitive triangle that are on one row from left to right, and then generate the next row right to left, and so on. Such a folded path ensures that an average distance from a generated pixel to recently previously generated pixels is relatively small.

Generating pixels that are near recently generated pixels is important when recent groups of pixels and/or their corresponding texture values are kept in memories of a limited size. The boustrophedonic sequence more often finds the pixels or texture values already loaded into such memories, and therefore repeating the memory load occurs less often.

In one embodiment, at least one boundary may be used which divides the primitive into a plurality of portions. In operation, the polygon-defining points may be moved in each of the portions separately. Further, the points may be moved through an entirety of a first one of the portions before being moved in a second one of the portions. As an option, the boundary spacing may be defined based on a memory size and the size of the smallest texture element, a "texel."

Rasterization with boundaries is another constraint on the pixel generation sequence that is complementary to boustrophedonic rasterization. If a primitive is very wide, the pixels of a single row may not fit in a limited-size memory. Rasterization with boundaries divides the triangle into limited-width rows (or columns), and generates all the pixels within such a portion before moving on to the next portion.

These and other advantages of the present invention will become apparent upon reading the following detailed description and studying the various figures of the drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other aspects and advantages are better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

FIG. 1 illustrates a prior art method of rasterization;

FIG. 1A illustrates the association of a set of edge variables with a particular pixel in accordance with a prior art method of rasterization;

FIG. 1B is a flow diagram illustrating the various components of one embodiment of the present invention implemented on a single semiconductor platform;

FIG. 2 is a schematic diagram of a vertex attribute buffer (VAB) in accordance with one embodiment of the present invention;

FIG. 2A is a chart illustrating the various commands that may be received by VAB in accordance with one embodiment of the present invention;

FIG. 2B is a flow chart illustrating a method of loading and draining vertex attributes to and from VAB in accordance with one embodiment of the present invention;

FIG. 2C is a schematic diagram illustrating the architecture of the present invention employed to implement the operations of FIG. 2B;

FIG. 3 illustrates the mode bits associated with VAB in accordance with one embodiment of the present invention;

FIG. 4 illustrates the transform module of the present invention;

FIG. 4A is a flow chart illustrating a method of running multiple execution threads in accordance with one embodiment of the present invention;

FIG. 4B is a flow diagram illustrating a manner in which the method of FIG. 4A is carried out in accordance with one embodiment of the present invention;

FIG. 5 illustrates the functional units of the transform module of FIG. 4 in accordance with one embodiment of the present invention;

FIG. 6 is a schematic diagram of the multiplication logic unit (MLU) of the transform module of FIG. 5;

FIG. 7 is a schematic diagram of the arithmetic logic unit (ALU) of the transform module of FIG. 5;

FIG. 8 is a schematic diagram of the register file of the transform module of FIG. 5;

FIG. 9 is a schematic diagram of the inverse logic unit (ILU) of the transform module of FIG. 5;

FIG. 10 is a chart of the output addresses of output converter of the transform module of FIG. 5 in accordance with one embodiment of the present invention;

FIG. 11 is an illustration of the micro-code organization of the transform module of FIG. 5 in accordance with one embodiment of the present invention;

FIG. 12 is a schematic diagram of the sequencer of the transform module of FIG. 5 in accordance with one embodiment of the present invention;



FIG. 13 is a flowchart delineating the various operations associated with use of the sequencer of the transform module of FIG. 12;

FIG. 14 is a flow diagram delineating the operation of the sequencing component of the sequencer of the transform module of FIG. 12;

FIG. 14A is a flow diagram illustrating the components of the present invention employed for handling scalar and vector components during graphics-processing;

FIG. 14B is a flow diagram illustrating one possible combination 1451 of the functional components of the present invention shown in FIG. 14A which corresponds to the transform module of FIG. 5;

FIG. 14C is a flow diagram illustrating another possible combination 1453 of the functional components of the present invention shown in FIG. 14A;

FIG. 14D illustrates a method implemented by the transform module of FIG. 12 for performing a blending operation during graphics-processing in accordance with one embodiment of the present invention;

FIG. 15 is a schematic diagram of the lighting module of one embodiment of the present invention;

FIG. 16 is a schematic diagram showing the functional units of the lighting module of FIG. 15 in accordance with one embodiment of the present invention;

FIG. 17 is a schematic diagram of the multiplication logic unit (MLU) of the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 18 is a schematic diagram of the arithmetic logic unit (ALU) of the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 19 is a schematic diagram of the register unit of the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 20 is a schematic diagram of the lighting logic unit (LLU) of the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 21 is an illustration of the flag register associated with the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 22 is an illustration of the micro-code fields associated with the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 23 is a schematic diagram of the sequencer associated with the lighting module of FIG. 16 in accordance with one embodiment of the present invention;

FIG. 24 is a flowchart delineating the manner in which the sequencers of the transform and lighting modules are capable of controlling the input and output of the associated buffers in accordance with one embodiment of the present invention;

FIG. 25 is a diagram illustrating the manner in which the sequencers of the transform and lighting modules are capable of controlling the input and output of the associated buffers in accordance with the method of FIG. 24;

FIG. 25B is a schematic diagram of the various modules of the rasterizer of FIG. 1B;

FIG. 26 illustrates a schematic of the set-up module of the rasterization module of the present invention;

FIG. 26A is an illustration showing the various parameters calculated by the set-up module of the rasterizer of FIG. 26;

FIG. 27 is a flowchart illustrating a method of the present invention associated with the set-up and traversal modules of the rasterizer component shown in FIG. 26;

FIG. 27A illustrates sense points that enclose a convex region that is moved to identify an area in a primitive in accordance with one embodiment of the present invention;

FIG. 28 is a flowchart illustrating a process of the present invention associated with the process row operation 2706 of FIG. 27;

FIG. 28A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive;

FIG. 28B illustrates another example of the sequence in which the convex region of the present invention is moved about the primitive;

FIG. 29 is a flowchart illustrating an alternate boustrophedonic process of the present invention associated with the process row operation 2706 of FIG. 27;

FIG. 29A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive in accordance with the boustrophedonic process of FIG. 29;

FIG. 30 is a flowchart illustrating an alternate boustrophedonic process using boundaries;

FIG. 31 is a flowchart showing the process associated with operation 3006 of FIG. 30;

FIG. 31A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive in accordance with the boundary-based boustrophedonic process of FIGS. 30 and 31;

FIG. 32 is a flowchart showing the process associated with operation 2702 of FIG. 27;

FIG. 32A is an illustration showing which area is drawn if no negative W-values are calculated in the process of FIG. 32;

FIG. 32B is an illustration showing which area is drawn if only one negative W-value is calculated in the process of FIG. 32; and

FIG. 32C is an illustration showing which area is drawn if only two negative W-values are calculated in the process of FIG. 32.

#### DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIGS. 1 and 1A show the prior art. FIGS. 1B-32C show a graphics pipeline system of the present invention.

FIG. 1B is a flow diagram illustrating the various components of one embodiment of the present invention. As shown, the present invention is divided into four main modules including a vertex attribute buffer (VAB) 50, a transform module 52, a lighting module 54, and a rasterization module 56 with a set-up module 57. In one embodiment, each of the foregoing modules is situated on a single semiconductor platform in a manner that will be described hereinafter in greater detail. In the present description, the single semiconductor platform may refer to a sole unitary semiconductor-based integrated circuit or chip.

The VAB 50 is included for gathering and maintaining a plurality of vertex attribute states such as position, normal, colors, texture coordinates, etc. Completed vertices are processed by the transform module 52 and then sent to the lighting module 54. The transform module 52 generates vectors for the lighting module 54 to light. The output of the lighting module 54 is screen space data suitable for the set-up module which, in turn, sets up primitives. Thereafter, rasterization module 56 carries out rasterization of the primitives. It should be noted that the transform and lighting



modules **52** and **54** might only stall on the command level such that a command is always finished once started.

In one embodiment, the present invention includes a hardware implementation that at least partially employs Open Graphics Library (OpenGL®) and D3D™ transform and lighting pipelines. OpenGL® is the computer industry's standard application program interface (API) for defining 2-D and 3-D graphic images. With OpenGL®, an application can create the same effects in any operating system using any OpenGL®-adhering graphics adapter. OpenGL® specifies a set of commands or immediately executed functions. Each command directs a drawing action or causes special effects.

FIG. 2 is a schematic diagram of VAB **50** in accordance with one embodiment of the present invention. As shown, VAB **50** passes command bits **200** while storing data bits **204** representative of attributes of a vertex and mode bits **202**. In use VAB **50** receives the data bits **204** of vertices and drains the same.

The VAB **50** is adapted for receiving and storing a plurality of possible vertex attribute states via the data bits **204**. In use after such data bits **204**, or vertex data, is received and stored in VAB **50**, the vertex data is outputted from VAB **50** to a graphics-processing module, namely the transform module **52**. Further, the command bits **200** are passed by VAB **50** for determining a manner in which the vertex data is inputted to VAB **50** in addition to other processing which will be described in greater detail with reference to FIG. 2A. Such command bits **200** are received from a command bit source such as a microcontroller, CPU, data source or any other type of source which is capable of generating command bits **200**.

Further, mode bits **202** are passed which are indicative of the status of a plurality of modes of process operations. As such, mode bits **202** are adapted for determining a manner in which the vertex data is processed in the subsequent graphics-processing modules. Such mode bits **202** are received from a command bit source such as a microcontroller, CPU, data source or any other type of source which is capable of generating mode bits **202**.

It should be noted that the various functions associated with VAB **50** may be governed by way of dedicated hardware, software or any other type of logic. In various embodiments, 64, 128, 256 or any other number of mode bits **202** may be employed.

The VAB **50** also functions as a gathering point for the 64 bit data that needs to be converted into a 128-bit format. The VAB **50** input is 64 bits/cycle and the output is 128 bits/cycle. In other embodiments, VAB **50** may function as a gathering point for 128-bit data, and VAB **50** input may be 128 bits/cycle or any other combination. The VAB **50** further has reserved slots for a plurality of vertex attributes that are all IEEE 32 bit floats. The number of such slots may vary per the desires of the user. Table 1 illustrates exemplary vertex attributes employed by the present invention.

TABLE 1

|                        |
|------------------------|
| Position: x,y,z,w      |
| Diffuse Color: r,g,b,a |
| Specular Color: r,g,b  |
| Fog: f                 |
| Texture0: s,t,r,q      |
| Texture1: s,t,r,q      |
| Normal: nx,ny,nz       |
| Skin Weight: w         |

During operation, VAB **50** may operate assuming that the x,y data pair is written before the z,w data pair since this allows for defaulting the z,w pair to (0.0,1.0) at the time of the x,y write. This may be important for default components in OpenGL™ and D3D™. It should be noted that the position, texture0 and texture1 slots default the third and fourth components to (0.0,1.0). Further, the diffuse color slot defaults the fourth component to (1.0) and the texture slots default the second component to (0.0).

The VAB **50** includes still another slot **205** used for assembling the data bits **204** that may be passed into or through the transform and lighting module **52** and **54**, respectively, without disturbing the data bits **204**. The data bits **204** in the slot **205** can be in a floating point or integer format. As mentioned earlier, the data bits **204** of each vertex has an associated set of mode bits **202** representative of the modes affecting the processing of the data bits **204**. These mode bits **202** are passed with the data bits **204** through the transform and lighting modules **52** and **54**, respectively, for purposes that will be set forth hereinafter in greater detail.

In one embodiment, there may be 18 valid VAB, transform, and lighting commands received by VAB **50**. FIG. 2A is a chart illustrating the various commands that may be received by VAB **50** in accordance with one embodiment of the present invention. It should be understood that all load and read context commands, and the passthrough command shown in the chart of FIG. 2A transfer one data word of up to 128 bits or any other size.

Each command of FIG. 2A may contain control information dictating whether each set of data bits **204** is to be written into a high double word or low double word of one VAB address. In addition, a 2-bit write mask may be employed for providing control to the word level. Further, there may be a launch bit that informs VAB controller that all of the data bits **204** are present for a current command to be executed.

Each command has an associated stall field that allows a look-up to find information on whether the command is a read command in that it reads context memory or is a write command in that it writes context memory. By using the stall field of currently executing commands, the new command may be either held off in case of conflict or allowed to proceed.

In operation, VAB **50** can accept one input data word up to 128 bits (or any other size) per cycle and output one data word up to 128 bits (or any other size) per cycle. For the load commands, this means that it may take two cycles to load the data into VAB **50** to create a 128-bit quad-word and one cycle to drain it. For the scalar memories in the lighting module **54**, it is not necessary to accumulate a full quad-word, and these can be loaded in one cycle/address. For one vertex, it can take up to 14 cycles to load the 7 VAB slots while it only takes 7 cycles to drain them. It should be noted, however, that it is only necessary to update the vertex state that changes between executing vertex commands. This means that, in one case, the vertex position may be updated taking 2 cycles, while the draining of the vertex data takes 7 cycles. It should be noted that only 1 cycle may be required in the case of the x,y position.

FIG. 2B is a flow chart illustrating one method of loading and draining vertex attributes to and from VAB **50** during graphics-processing. Initially, in operation **210**, at least one set of vertex attributes is received in VAB **50** for being processed. As mentioned earlier, each set of vertex attributes may be unique, and correspond to a single vertex.

In use the vertex attributes are stored in VAB **50** upon the receipt thereof in operation **212**. Further, each set of stored



vertex attributes is transferred to a corresponding one of a plurality of input buffers of the transform module 52. The received set of vertex attributes is also monitored in order to determine whether a received vertex attribute has a corresponding vertex attribute of a different set currently stored in VAB 50, as indicated in operation 216.

Upon it being determined that a stored vertex attribute corresponds to the received vertex attribute in decision 217, the stored vertex attribute is outputted to the corresponding input buffer of the transform module 52 out of order. See operation 218. Immediately upon the stored vertex attribute being outputted, the corresponding incoming vertex attribute may take its place in VAB 50. If no correspondence is found, however, each set of the stored vertex attributes may be transferred to the corresponding input buffer of the transform module 52 in accordance with a regular predetermined sequence. Note operation 219.

It should be noted that the stored vertex attribute might not be transferred in the aforementioned manner if it has an associated launch command. Further, in order for the foregoing method to work properly, the bandwidth of an output of VAB 50 must be at least the bandwidth of an input of VAB 50.

FIG. 2C is a schematic diagram illustrating the architecture of the present invention employed to implement the operations of FIG. 2B. As shown, VAB 50 has a write data terminal WD, a read data terminal RD, a write address terminal WA, and a read address RA terminal. The read data terminal is coupled to a first clock-controlled buffer 230 for outputting the data bits 204 from VAB 50.

Also included is a first multiplexer 232 having an output coupled to the read address terminal of VAB 50 and a second clock-controlled buffer 234. A first input of the first multiplexer 232 is coupled to the write address terminal of VAB 50 while a second input of the first multiplexer 232 is coupled to an output of a second multiplexer 236. A logic module 238 is coupled between the first and second multiplexers 232 and 236, the write address terminal of VAB 50, and an output of the second clock-controlled buffer 234.

In use the logic module 238 serves to determine whether an incoming vertex attribute is pending to drain in VAB 50. In one embodiment, this determination may be facilitated by monitoring a bit register that indicates whether a vertex attribute is pending or not. If it is determined that the incoming vertex attribute does have a match currently in VAB 50, the logic module 238 controls the first multiplexer 232 in order to drain the matching vertex attribute so that the incoming vertex attribute may be immediately stored in its place. On the other hand, if it is determined that the incoming vertex attribute does not have a match currently in VAB 50, the logic module 238 controls the first multiplexer 232 such that VAB 50 is drained and the incoming vertex attribute is loaded sequentially or in some other predetermined order, per the input of the second multiplexer 236 which may be updated by the logic module 238.

As a result, there is no requirement for VAB 50 to drain multiple vertex attributes before a new incoming vertex attribute may be loaded. The pending vertex attribute forces out the corresponding VAB counterpart if possible, thus allowing it to proceed. As a result, VAB 50 can drain in an arbitrary order. Without this capability, it would take 7 cycles to drain VAB 50 and possibly 14 more cycles to load it. By overlapping the loading and draining, higher performance is achieved. It should be noted that this is only possible if an input buffer is empty and VAB 50 can drain into input buffers of the transform module 52.

FIG. 3 illustrates the mode bits associated with VAB 50 in accordance with one embodiment of the present invention. The transform/light mode information is stored in a register via mode bits 202. Mode bits 202 are used to drive the sequencers of the transform module 52 and lighting module 54 in a manner that will become apparent hereinafter. Each vertex has associated mode bits 202 that may be unique, and can therefore execute a specifically tailored program sequence. While, mode bits 202 may generally map directly to the graphics API, some of them may be derived.

In one embodiment, the active light bits (LIS) of FIG. 3 may be contiguous. Further, the pass-through bit (VPAS) is unique in that when it is turned on, the vertex data is passed through with scale and bias, and no transforms or lighting is done. Possible mode bits 202 used when VPAS is true are the texture divide bits (TDV0,1), and foggen bits (used to extract fog value in D3D™). VPAS is thus used for pre-transformed data, and TDV0,1 are used to deal with a cylindrical wrap mode in the context of D3D™.

FIG. 4 illustrates the transform module of one embodiment of the present invention. As shown, the transform module 52 is connected to VAB 50 by way of 6 input buffers 400. In one embodiment, each input buffer 400 might be 7\*128b in size. The 6 input buffers 400 each is capable of storing 7 quad words. Such input buffers 400 follow the same layout as VAB 50, except that the pass data is overlapped with the position data.

In one embodiment, a bit might be designated for each attribute of each input buffer 400 to indicate whether data has changed since the previous instance that the input buffer 400 was loaded. By this design, each input buffer 400 might be loaded only with changed data.

The transform module 52 is further connected to 6 output vertex buffers 402 in the lighting module 54. The output buffers include a first buffer 404, a second buffer 406, and a third buffer 408. As will become apparent hereinafter, the contents, i.e. position, texture coordinate data, etc., of the third buffer 408 are not used in the lighting module 54. The first buffer 404 and second buffer 406 are both, however, used for inputting lighting and color data to the lighting module 54. Two buffers are employed since the lighting module is adapted to handle two read inputs. It should be noted that the data might be arranged so as to avoid any problems with read conflicts, etc.

Further coupled to the transform module 52 is context memory 410 and micro-code ROM memory 412. The transform module 52 serves to convert object space vertex data into screen space, and to generate any vectors required by the lighting module 54. The transform module 52 also does processes skinning and texture coordinates. In one embodiment, the transform module 52 might be a 128-bit design processing 4 floats in parallel, and might be optimized for doing 4 term dot products.

FIG. 4A is a flow chart illustrating a method of executing multiple threads in the transform module 52 in accordance with one embodiment of the present invention. In operation, the transform module 52 is capable of processing 3 vertices in parallel via interleaving. To this end, 3 commands can be simultaneously executed in parallel unless there are stall conditions between the commands such as writing and subsequently reading from the context memory 410. The 3 execution threads are independent of each other and can be any command since all vertices contain unique corresponding mode bits 202.

As shown in FIG. 4A, the method of executing multiple threads includes determining a current thread to be executed



in operation 420. This determination might be made by identifying a number of cycles that a graphics-processing module requires for completion of an operation, and tracking the cycles. By tracking the cycles, each thread can be assigned to a cycle, thus allowing determination of the current thread based on the current cycle. It should be noted, however, that such determination might be made in any desired manner that is deemed effective.

Next, in operation 422, an instruction associated with a thread to be executed during a current cycle is retrieved using a corresponding program counter number. Thereafter, the instruction is executed on the graphics-processing module in operation 424.

In one example of use, the instant method includes first accessing a first instruction, or code segment, per a first program counter. As mentioned earlier, such program counter is associated with a first execution thread. Next, the first code segment is executed in the graphics-processing module. As will soon become apparent, such graphics-processing module might take the form of an adder, a multiplier, or any other functional unit or combination thereof.

Since the graphics-processing module requires more than one clock cycle to complete the execution, a second code segment might be accessed per a second program counter immediately one clock cycle after the execution of the first code segment. The second program counter is associated with a second execution thread, wherein each of the execution threads process a unique vertex.

To this end, the second code segment might begin execution in the graphics-processing module prior to the completion of the execution of the first code segment in the graphics-processing module. In use the graphics-processing module requires a predetermined number of cycles for every thread to generate an output. Thus, the various steps of the present example might be repeated for every predetermined number of cycles.

This technique offers numerous advantages over the prior art. Of course, the functional units of the present invention are used more efficiently. Further, the governing code might be written more efficiently when the multiple threading scheme is assumed to be used.

For example, in the case where the graphics-processing module includes a multiplier that requires three clock cycles to output an answer, it would be necessary to include two no operation commands between subsequent operations such as  $a=b*c$  and  $d=e*a$ , since "a" would not be available until after the three clock cycles. In the present embodiment, however, the code might simply call  $d=e*a$  immediately subsequent  $a=b*c$ , because it can be assumed that such code will be executed as one of three execution threads that are called once every three clock cycles.

FIG. 4B is a flow diagram illustrating a manner in which the method of FIG. 4A is carried out. As shown, each execution thread has an associated program counter 450 that is used to access instructions, or code segments, in instruction memory 452. Such instructions might then be used to operate a graphics-processing module such as an adder 456, a multiplier 454, and/or an inverse logic unit or register 459.

In order to accommodate a situation where at least two of the foregoing processing modules are used in tandem, at least one code segment delay 457 is employed between the graphics-processing modules. In the case where a three-thread framework is employed, a three-clock cycle code segment delay 457 is used. In one embodiment, the code segment delay 457 is used when a multiplication instruction

is followed by an addition instruction. In such case, the addition instruction is not executed until three clock cycles after the execution of the multiplication instruction in order to ensure that time has elapsed which is sufficient for the multiplier 456 to generate an output.

After the execution of each instruction, the program counter 450 of the current execution thread is updated and the program counter of the next execution thread is called by module 458 in a round robin sequence to access an associated instruction. It should be noted that the program counters might be used in any fashion including, but not limited to incrementing, jumping, calling and returning, performing a table jump, and/or dispatching. Dispatching refers to determining a starting point of code segment execution based on a received parameter. Further, it important to understand that the principles associated with the present multiple thread execution framework might also be applied to the lighting module 54 of the graphics-processing pipeline of the present invention.

In the case where a three-thread framework is employed, each thread is allocated one input buffer and one output buffer at any one time. This allows loading of three more commands with data while processing three commands. The input buffers and output buffers are assigned in a round robin sequence in a manner that will be discussed later with reference to FIGS. 27 and 28.

The execution threads are thus temporally and functionally interleaved. This means that each function unit is pipelined into three stages and each thread occupies one stage at any one time. In one embodiment, the three-threads might be set to always execute in the same sequence, i.e. zero then one then three. Conceptually, the threads enter a function unit at  $t=\text{clock modulo three}$ . Once a function unit starts work, it takes three cycles to deliver the result (except the ILU that takes six), at which time the same thread is again active.

FIG. 5 illustrates the functional units of the transform module 52 of FIG. 4 in accordance with one embodiment of the present invention. As shown, included are input buffers 400 that are adapted for being coupled to VAB 50 for receiving vertex data therefrom.

A memory logic unit (MLU) 500 has a first input coupled to an output of input buffers 400. As an option, the output of MLU 500 might have a feedback loop 502 coupled to the first input thereof.

Also provided is an arithmetic logic unit (ALU) 504 having a first input coupled to an output of MLU 500. The output of ALU 504 further has a feedback loop 506 connected to the second input thereof. Such feedback loop 502 may further have a delay 508 coupled thereto. Coupled to an output of ALU 504 is an input of a register unit 510. It should be noted that the output of register unit 510 is coupled to the first and second inputs of MLU 500.

An inverse logic unit (ILU) 512 is provided including an input coupled to the output of ALU 504 for performing an inverse or an inverse square root operation. In an alternate embodiment, ILU 512 might include an input coupled to the output of register unit 510.

Further included is a conversion, or smearing, module 514 coupled between an output of ILU 512 and a second input of MLU 500. In use the conversion module 514 serves to convert scalar vertex data to vector vertex data. This is accomplished by multiplying the scalar data by a vector so that the vector operators such as the multiplier and/or adder may process it. For example, a scalar A, after conversion, may become a vector (A,A,A,A). In an alternate



embodiment, the smearing module 514 might be incorporated into the multiplexers associated with MLU 500, or any other component of the present invention. As an option, a register 516 might be coupled between the output of ILU 512 and an input of the conversion unit 514. Further, such register 516 might be threaded.

Memory 410 is coupled to the second input of MLU 500 and the output of ALU 504. In particular, memory 410 has a read terminal coupled to the second input of MLU 500. Further, memory 410 has a write terminal coupled to the output of ALU 504.

The memory 410 has stored therein a plurality of constants and variables for being used in conjunction with the input buffer 400, MLU 500, ALU 504, register unit 510, ILU 512, and the conversion module 514 for processing the vertex data. Such processing might include transforming object space vertex data into screen space vertex data, generating vectors, etc.

Finally, an output converter 518 is coupled to the output of ALU 504. The output converter 518 serves for being coupled to a lighting module 54 via output buffers 402 to output the processed vertex data thereto. All data paths except for the ILU might be designed to be 128 bits wide or other data path widths may be used.

FIG. 6 is a schematic diagram of MLU 500 of the transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. As shown, MLU 500 of the transform module 52 includes four multipliers 600 that are coupled in parallel.

MLU 500 of transform module 52 is capable of multiplying two four component vectors in three different ways, or pass one four component vector. MLU 500 is capable of performing multiple operations. Table 2 illustrates such operations associated with MLU 500 of transform module 52.

TABLE 2

|           |  |
|-----------|--|
| CMLU_MULT | $o[0] = a[0]*b[0], o[1] = a[1]*b[1], o[2] = a[2]*b[2], o[3] = a[3]*b[3]$ |
| CMLU_MULA | $o[0] = a[0]*b[0], o[1] = a[1]*b[1], o[2] = a[2]*b[2], o[3] = a[3]$      |
| CMLU_MULB | $o[0] = a[0]*b[0], o[1] = a[1]*b[1], o[2] = a[2]*b[2], o[3] = b[3]$      |
| CMLU_PASA | $o[0] = a[0], o[1] = a[1], o[2] = a[2], o[3] = a[3]$                     |
| CMLU_PASB | $o[0] = b[0], o[1] = b[1], o[2] = b[2], o[3] = b[3]$                     |

Possible A and B inputs are shown in Table 3.

TABLE 3

|      |              |
|------|--------------|
| MA_M | MLU          |
| MA_V | Input Buffer |

TABLE 3-continued

|      |                        |
|------|------------------------|
| MA_R | RLU (shared with MB_R) |
| MB_I | ILU                    |
| MB_C | Context Memory         |
| MB_R | RLU (shared with MA_R) |

Table 4 illustrates a vector rotate option capable of being used for cross products.

TABLE 4

|         |  |
|---------|--|
| MR_NONE | No change                                      |
| MR_ALBR | Rotate A[XYZ] vector left, B[XYZ] vector right |
| MR_ARBL | Rotate A[XYZ] vector right, B[XYZ] vector left |

FIG. 7 is a schematic diagram of ALU 504 of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. As shown, ALU 504 of transform module 52 includes three adders 700 coupled in parallel/series. In use ALU 504 of transform module 52 can add two three component vectors, pass one four component vector, or smear a vector component across the output. Table 5 illustrates various operations of which ALU 504 of transform module 52 is capable.

TABLE 5

|            |   |
|------------|---|
| CALU_ADDA  | $o[0] = a[0]+b[0], o[1] = a[1]+b[1], o[2] = a[2]+b[2], o[3] = a[3]$ |
| CALU_ADDB  | $o[0] = a[0]+b[0], o[1] = a[1]+b[1], o[2] = a[2]+b[2], o[3] = b[3]$ |
| CALU_SUM3B | $o[0123] = b[0] + b[1] + b[2]$                                      |
| CALU_SUM4B | $o[0123] = b[0] + b[1] + b[2] + b[3]$                               |
| CALU_SMRB0 | $o[0123] = b[0]$  |
| CALU_SMRB1 | $o[0123] = b[1]$  |
| CALU_SMRB2 | $o[0123] = b[2]$  |
| CALU_SMRB3 | $o[0123] = b[3]$  |
| CALU_PASA  | $o[0] = a[0], o[1] = a[1], o[2] = a[2], o[3] = a[3]$                |
| CALU_PASB  | $o[0] = b[0], o[1] = b[1], o[2] = b[2], o[3] = b[3]$                |



Table 6 illustrates the A and B inputs of ALU 504 of transform module 52.

TABLE 6

|      |                             |
|------|-----------------------------|
| AA_A | ALU (one instruction delay) |
| AA_C | Context Memory              |
| AB_M | MLU                         |

It is also possible to modify the sign bits of the A and B input by effecting no change, negation of B, negation of A, absolute value A,B. It should be noted that when ALU 504 outputs scalar vertex data, this scalar vertex data is smeared across the output in the sense that each output represents the scalar vertex data. The pass control signals of MLU 500 and ALU 504 are each capable of disabling all special value handling during operation.

FIG. 8 is a schematic diagram of the vector register file 510 of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. As shown, the vector register file 510 includes four sets of registers 800 each having an output connected to a first input of a corresponding multiplexer 802 and an input coupled to a second input of the corresponding multiplexer 802.

In one embodiment of the present invention, the vector register file 510 is threaded. That is, there are three copies of the vector register file 510 and each thread has its own copy. In one embodiment, each copy contains eight registers, each of which might be 128 bits in size and store four floats. The vector register file 510 is written from ALU 504 and the output is fed back to MLU 500. The vector register file 510 has one write and one read per cycle.

In operation, it is also possible to individually mask a write operation to each register component. The vector register file 510 exhibits zero latency when the write address is the same as the read address due to a bypass path 511 from the input to the output. In this case, unmasked components would be taken from the registers and masked components would be bypassed. The vector register file 510 is thus very useful for building up vectors component by component, or for changing the order of vector components in conjunction with the ALU SMR operations (See Table 5). Temporary results might be also stored in the vector register file 510.

FIG. 9 is a schematic diagram of ILU 512 of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. As shown, ILU 512 of transform module 52 is capable of generating a floating-point reciprocal (1/D) and a reciprocal square root (1/D<sup>1/2</sup>). To carry out such operations, either one of two iterative processes might be executed on a mantissa. Such processes might be executed with any desired dedicated hardware, and are shown below:

| Reciprocal (1/D)                           | Reciprocal Square-root (1/D <sup>1/2</sup> ) |
|--|--|
| $x_{n+1} = x_n(2-x_n*D)$                   | $x_{n+1} = (1/2)*x_n(3-x_n^2*D)$             |
| 1) table look up for $x_n$ (seed)          | table look up for $x_n$ (seed)               |
| $x_n$                                      | $x_n * x_n$                                  |
| 2) 1 <sup>st</sup> iteration: multiply-add | 1 <sup>st</sup> iteration: multiply-add      |
| $2-x_n*D$                                  | $3-x_n^2*D$                                  |
| 3) 1 <sup>st</sup> iteration: multiply     | 1 <sup>st</sup> iteration: multiply          |
| $x_n(2-x_n*D)$                             | $(1/2)*x_n(3-x_n^2*D)$                       |
| 4) 2 <sup>nd</sup> iteration: no-op        | 2 <sup>nd</sup> iteration: square            |
| pass $x_{n+1}$                             | $x_{n+1}^2$                                  |
| 5) 2 <sup>nd</sup> iteration: multiply-add | 2 <sup>nd</sup> iteration: multiply-add      |
| $2-x_{n+1}*D$                              | $3-x_{n+1}^2*D$                              |

-continued

|       | Reciprocal (1/D)                    | Reciprocal Square-root (1/D <sup>1/2</sup> ) |
|-------|-------------------------------------|--|
| 5) 6) | 2 <sup>nd</sup> iteration: multiply | 2 <sup>nd</sup> iteration: multiply          |
|       | $x_{n+1} (2-x_{n+1}*D)$             | $(1/2)*x_{n+1}(3-x_{n+1}^2*D)$               |

As shown, the two processes are similar, affording a straightforward design. It should be noted that the iterations might be repeated until a threshold precision is met.

In operation, ILU 512 performs two basic operations including an inverse operation and inverse square root operation. Unlike the other units, it requires six cycles to generate the output. The input is a scalar, and so is the output. As set forth earlier, the threaded holding register 516 at ILU 512 output is relied upon to latch the result until the next time a valid result is generated. Further, the scalar output is smeared into a vector before being fed into MLU 500. The inverse unit 512 uses look-up tables and a two pass Newton-Raphson iteration to generate IEEE (Institute of Electrical and Electronics Engineers) outputs accurate to within about 22 mantissa bits. Table 7 illustrates the various operations that might be performed by ILU 512 of transform module 52.

TABLE 7

|           |                              |
|-----------|------------------------------|
| CILU_INV  | o = 1.0/a                    |
| CILU_ISQ  | o = 1.0/sqrt(a)              |
| CILU_CINV | o = 1.0/a (with range clamp) |
| CILU_NOP  | no output                    |

The foregoing range clamp inversion operation of Table 7 might be used to allow clipping operations to be handled by rasterization module 56. Coordinates are transformed directly into screen space that can result in problems when the homogeneous clip space w is near 0.0. To avoid multiplying by 1.0/0.0 in the perspective divide, the 1/w calculation is clamped to a minimum and a maximum exponent.

In use the context memory 410 as shown in FIG. 5 reads and writes only using quad-words. The memory can be read by MLU 500 or ALU 504 each cycle, and can be written by ALU 504. Only one memory read is allowed per cycle. If a read is necessary, it is done at the start of an instruction and then pipelined down to ALU 504 three cycles later. Context memory 410 need not necessarily be threaded.

FIG. 10 is a chart of the output addresses of output converter 518 of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. The output converter 518 is responsible for directing the outputs to proper destinations, changing the bit precision of data, and some data swizzling to increase performance. All data destined for lighting module 54 is rounded to a 22 bit floating point format organized as S1E8M13 (one sign, eight exponent, 13 mantissa bits). The destination buffers 402 as shown in FIG. 4 in lighting module 54 are threaded.

Data swizzling is useful when generating vectors. Such technique allows the generation of a distance vector (1,d,d\*d) without penalty when producing a vector.

The distance vector is used for fog, point parameter and light attenuation. This is done with an eye vector and light direction vectors. Table 8 illustrates the various operations associated with such vectors. It should be noted that, in the following table, squaring the vector refers to  $d^2 = \text{dot}[(x,y,z), (x,y,z)]$ , and storing  $d^2$  in the w component of (x,y,z).



TABLE 8

|                                 |                 |   |
|---------------------------------|-----------------|---|
| 1. Square the vector            | (x,y,z,d*d)     | (output d*d to VBUF, 1.0 to VBUF)       |
| 2. Generate inverse sqrt of d*d | (1/d)           |   |
| 3. Normalize vector             | (x/d,y/d,z/d,d) | (output x/d,y/d,z/d to WBUF, d to VBUF) |

It should be noted that the math carried out in the present invention might not always be IEEE compliant. For example, it might be assumed that "0" multiplied by any number renders "0." This is particularly beneficial when dealing with the equations such as  $d=d^2*1/(d^2)^{1/2}$ , where  $d=0$ . Without making the foregoing assumption, such equation would afford an error, thus causing problems in making related computations.

FIG. 11 is an illustration of the micro-code organization of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. The transform module micro-code might be arranged into 15 fields making up a total width of 44 bits. Fields might be delayed to match the data flow of the units. MLU 500 operations are executed at a delay of zero, ALU operations are executed at a delay of one, and RLU, output operations are executed at a delay of two. Each delay is equivalent to three cycles.

FIG. 12 is a schematic diagram of sequencer 1200 of transform module 52 of FIG. 5 in accordance with one embodiment of the present invention. As shown in FIG. 12, sequencer 1200 of transform module 52 includes a buffer 1202 adapted for receiving the mode bits from VAB 50 that are indicative of the status of a plurality of modes of process operations.

Also included is memory 412 capable of storing code segments that each are adapted to carry out the process operations in accordance with the status of the modes. A sequencing module 1206 is coupled between memory 412 and a control vector module 1205 which is in turn coupled to buffer 1202 for identifying a plurality of addresses in memory 412 based on a control vector derived from mode bits 202. The sequencing module 1206 is further adapted for accessing the addresses in memory 412 for retrieving the code segments that might be used to operate transform module 52 to transfer data to an output buffer 1207.

FIG. 13 is a flowchart delineating the various operations associated with use of sequencer 1200 of transform module 52 of FIG. 12. As shown, sequencer 1200 is adapted for sequencing graphics-processing in a transform or lighting operation. In operation 1320, mode bits 202 are first received which are indicative of the status of a plurality of modes of process operations. In one embodiment, mode bits 202 might be received from a software driver.

Then, in operation 1322, pluralities of addresses are then identified in memory based on mode bits 202. Such addresses are then accessed in the memory in operation 1324 for retrieving code segments that each are adapted to carry out the process operations in accordance with the status of the modes. The code segments are subsequently executed with a transform or lighting module for processing vertex data. Note operation 1326.

FIG. 14 is a flow diagram delineating the operation of the sequencing module 1206 of sequencer 1200 of transform module 52 of FIG. 12. As shown, a plurality of mode registers 1430 each include a unique set of mode bits 202 which in turn correspond to a single vertex. It should be noted that mode registers 1430 are polled in a round robin sequence in order to allow the execution of multiple execution threads in the manner set forth earlier during reference to FIGS. 4A and 4B.

Once the current execution thread is selected, a corresponding group of mode bits 202 are decoded in operation 1432. Upon mode bits 202 being decoded in operation 1432, a control vector is afforded which includes a plurality of bits each of which indicate whether a particular code segment is to be accessed in ROM 1404 for processing the corresponding vertex data.

Upon determining whether a code segment should be accessed in ROM 1404 and executed, a pointer operation 1436 increments the current thread pointer to start the next execution thread to obtain a second group mode bits 202 to continue a similar operation. This might be continued for each of the threads in a round robin sequence.

Once the control vector has been formed for a particular group of mode bits 202, a priority encoder operation 1438 determines, or identifies, a next "1" or enabled, bit of the control vector. If such a bit is found, the priority encoder operation 1438 produces an address in ROM 1404 corresponding to the enabled bit of the control vector for execution purposes.

Upon returning to the initial group of mode bits 202 after handling the remaining threads, and after the mode bits have been decoded and the control vector is again available, a masking operation 1434 might be used to mask the previous "1", or enabled, bit that was identified earlier. This allows analysis of all remaining bits after mask operation 1434.

The foregoing process might be illustrated using the following tables. Table 9 shows a plurality of equations that might be executed on subject vertex data.

TABLE 9

|     |                    |   |
|-----|--------------------|---|
| R = | (a                 | ) |
| R = | (a + d*e           | ) |
| R = | (a + b*c + f)      |   |
| R = | (a + b*c + d*e     | ) |
| R = | 1.0/(a             | ) |
| R = | 1.0/(a + d*e       | ) |
| R = | 1.0/(a + b*c + f)  |   |
| R = | 1.0/(a + b*c + d*e | ) |

As shown, there are four possibilities of products that might be summed in addition to an inverse operation (a, b\*c, d\*e, f, and 1/x). Next, mode fields might be defined. Table 10 illustrates a pair of mode fields, mode.y and mode.z, each having assigned thereto a predetermined set of the operations of Table 9.

TABLE 10

|           |                      |
|-----------|----------------------|
| mode.y[4] | 0: R = a             |
|           | 1: R = a + d*e       |
|           | 2: R = a + b*c + f   |
|           | 3: R = a + b*c + d*e |
| mode.z[2] | 0: R = R             |
|           | 1: R = 1.0/R         |

Thereafter, each of the operations might be positioned in memory with an associated address. Table 11 illustrates a plurality of memory addresses each having an associated operation. Also shown is a set of control vector definitions.

TABLE 11

|         |             |
|---------|-------------|
| ROM[0]: | R = a       |
| ROM[1]: | R = R + b*c |
| ROM[2]: | R = R + d*e |
| ROM[3]: | R = R + f   |



TABLE 11-continued

---

|   |
|---|
| ROM[4]: R = 1.0/R                         |
| cv[0] = 1;                                |
| cv[1] = (mode.y==2    mode.y==3) ? 1 : 0; |
| cv[2] = (mode.y==1    mode.y==3) ? 1 : 0; |
| cv[3] = (mode.y==2) ? 1 : 0;              |
| cv[4] = (mode.z==1) ? 1 : 0;              |

---

Table 12 illustrates the execution of an example.

TABLE 12

---

|   |
|---|
| R = a+d*e corresponds to:                               |
| mode.y = 1;   |
| mode.z = 0;   |
| which in turn affords the following control vector:     |
| cv[0] = 1;  |
| cv[1] = 0;  |
| cv[2] = 1;  |
| cv[3] = 0;  |
| cv[4] = 0;  |
| execution   |
| first cycle:  |
| cv[0] is TRUE so execute ROM[0]                         |
| more TRUE values in control vector, so do not terminate |
| program   |
| second cycle:   |
| cv[1] is FALSE so keep looking                          |
| cv[2] is TRUE so execute ROM[2]                         |
| no more TRUE values in control vector, so terminate     |
| program   |

---

As such, sequencer 1200 of transform module 52 steps through a threaded control vector which is derived from threaded mode bits 202, and executes every ROM address whose corresponding control vector bit is set to "TRUE". The control vector has the same length as the ROM. The sequencer 1200 is capable of stepping through an arbitrary control vector at the rate of one "1", or enabled bit per a predetermined number of cycles. Commands that do not use mode bits 202 might be executed by on-the-fly micro-code generation due to the simplicity thereof.

By representing such statuses by way of a unique string of mode bits 202, it is unnecessary to execute a plurality of if-then clauses in the graphics-processing hardware to determine the statuses of the various operations. Improved performance is thereby afforded. Conceptually, it is as if the if clauses in a program language had been moved to sequencer 1200 which in turn instantly skips instructions with a "FALSE" condition, as indicated by mode bits 202.

As indicated earlier, code segments are stored in the ROM which are capable of handling the various statuses of the operations identified by the mode bits. In one embodiment a separate code segment might be retrieved for handling each operation indicated by the mode bits. In the alternative, a single comprehensive code segment might be written for handling each or some combinations of operations that are possible. It should be noted, however, that generating such large code segments for each combination of operations requires additional code space, and it therefore might be beneficial to modularize the code segments for only commonly used combinations of operations.

Since mode bits 202 do not change once the vertex commences execution, the control vector generation might only have to be done once per vertex before entering the sequencer. Exceptions to this might arise in some cases, however, such as lighting where operations might be repeated. When the last vertex instruction is found, an end of sequence (EOS) signal might be asserted. This in turn

might be used to change the status of the input and output buffers, and to allow the start of the next command in a manner that will be set forth during reference to FIGS. 28A and 28B. It should be noted that the EOS signal is pipeline delayed for release of the destination buffer similar to the manner in which the instructions are handled. See FIG. 4B.

FIG. 14A is a flow diagram illustrating the various functional components of the present invention employed for integrating the handling of scalar and vector vertex data during graphics-processing. As shown, one functional aspect 1440 includes inputting vector vertex data into a processing module, i.e. adder, multiplier, etc., for outputting vector vertex data. In another functional aspect 1442, vector vertex data is processed by a vector processing module, i.e. adder, multiplier, etc., which outputs scalar vertex data that is in turn converted, or smeared, again into vector vertex data.

In yet another functional aspect 1444, vector vertex data is masked, thereby converted to scalar vertex data, after which it is stored in memory, i.e. register logic unit, for the purpose of generating vector vertex data. In still yet another functional aspect 1446, scalar vertex data is extracted by a vector processing module, i.e. adder, multiplier, etc., which in turn is processed by a scalar processing module, i.e. inverse logic unit, which renders scalar vertex data. This scalar vertex data is converted again into vector vertex data.

FIG. 14B is a flow diagram illustrating one possible combination 1451 of the functional components of the present invention shown in FIG. 14A which corresponds to transform module 52 of FIG. 5. It should be noted that functional aspects 1444 and 1446 might have delays associated therewith in a manner similar to that set forth earlier during reference to FIG. 4B. FIG. 14C is a flow diagram illustrating yet another possible combination 1453 of the functional components of the present invention shown in FIG. 14A.

Multiplexers might accomplish the extraction of the scalar vertex data from the vector vertex data in the functional modules of FIGS. 14A-14C. Such multiplexers might also be responsible for any data swizzling that might be required before processing by the various functional modules. In one embodiment, the multiplexers might be capable of passing and rotating vector vertex data, and rely on other graphics-processing modules such as an ALU for other processing. In yet another embodiment, the multiplexers might be capable of arbitrarily rearranging attributes independently without penalty.

FIG. 14D illustrates a method in which the transform system is adapted for performing a blending, or skinning operation during graphics-processing in a graphics pipeline via a hardware implementation such as an application specific integrated circuit (ASIC). During processing in the pipeline, in operation 1470, a plurality of matrices, a plurality of weight values each corresponding with one of the matrices, and vertex data are received. It should be noted that an additional set of matrices might be required for normal vertex data.

Subsequently, in operation 1472, a sum of a plurality of products is then calculated with each product being calculated by the multiplication of the vertex data, one of the matrices and the weight corresponding to the matrix. Such sum of products is then outputted in operation 1474 for additional processing.

In summary, the following sum of products might be calculated:

$$v' = \sum w_i * M_i * v \text{ for } i=1 \dots x$$

Equation #1



where  $v$ =inputted vertex data

$w$ =weight value

$M$ =matrix

$x$ =number of matrices

$v'$ =vertex data for output to a processing module

$$n' = \sum w_i * I_i * n \text{ for } i=1 \dots x \quad \text{Equation \#2}$$

where  $n$ =inputted vertex data (normal vector)

$w$ =weight value

$I$ =inverted matrix (inverse transpose matrix)

$x$ =number of inverted matrices

$n'$ =vertex data for output to a processing module (normal vector)

$$v_s [O_x, O_y, O_z, \emptyset] + 1 / (v''_{wc}) * [(v''_x), (v''_y), (v''_z), 1] \quad \text{Equation \#3}$$

where

$v'' = C * v'$

$v'$ =sum of products from Equation #1

$C = [S_x, S_y, S_z, 1] * P$

$P$ =projection matrix

$v_s$ =screen vector for display purposes

$O$ =viewport offset

$S$ =viewport scale

It should be noted that there are many ways to represent the weights  $w_i$  set forth hereinabove. For example, in Equations #1 and #2 above, it might be said that  $i=1 \dots (x-1)$ , leaving  $w_x$  ( $w_i$  where  $i=x$ ) to be calculated by the equation  $1 - \sum w_i$ . By representing the weights  $w_i$  in this way, it is ensured that all of the weights  $w$  sum to 1.

In one embodiment, the matrices might include model view matrices ( $M$ ), and the sum of products ( $v'$ ) might be outputted for additional processing by a lighting operation. See Equation #1. This sum of products ( $v'$ ) might also be used to generate another sum of products ( $v_s$ ) for display purposes by using a composite matrix ( $C$ ). See Equation #3. Still yet, the matrices might include inverse matrices ( $I$ ) and the vertex data might include normal vector data ( $n$ ). In such case, the additional processing might include a lighting operation. See Equation #2.

FIG. 15 is a schematic diagram of lighting module 54 in accordance with one embodiment of the present invention. As shown, lighting module 54 includes buffers 402 to which transform module 52 outputs the vertex data. As shown, buffer 408 bypasses lighting module 54 by way of the pathway 1501. Further coupled to lighting module 54 is a context memory 1500 and micro-code ROM memory 1502.

The lighting module 54 is adapted for handling lighting in addition to fog and point parameters. In use lighting module 54 controls the buffer bypass pathway 1501, and calculates the diffuse, point size, and specular output colors as well as the fog value. It should be noted that lighting module 54 employs the same mode bits 202 as transform module 52.

The lighting module 54 further requires less precision with respect to transform module 52, and therefore processes 22 bit floating point values (1.8.13 format) organized in tri-words. Since the data of third buffer 408 is 128 bits, it utilizes bypass pathway 1501 around lighting module 54. The lighting module 54 is event driven and simultaneously executes three threads in a manner similar to transform module 52 as was set forth earlier with reference to FIGS. 4A and 4B. It should be noted that lighting module 54 might require command launch approval from an outside source.

FIG. 16 is a schematic diagram showing the functional units of lighting module 54 of FIG. 15 in accordance with

one embodiment of the present invention. As shown, included are input buffers 402 adapted for being coupled to a transform system for receiving vertex data therefrom. As set forth earlier, input buffers 402 include a first input buffer 404, a second input 406, and a third input buffer 408. An input of first buffer 404, second input buffer 406, and third input buffer 408 are coupled to an output of transform module 52. For bypass purposes, the output of third buffer 408 is coupled to the output of lighting module 54 via a delay 1608.

Further included is a MLU 1610 having a first input coupled to an output of first input buffer 404 and a second input coupled to an output of second input buffer 406. The output of MLU 1610 has a feedback loop 1612 coupled to the second input thereof. An arithmetic logic unit (ALU) 1614 has a first input coupled to an output of second input buffer 406. ALU 1614 further has a second input coupled to an output of MLU 1610. An output of ALU 1614 is coupled to the output of lighting module 54. It should be noted that the output of ALU 1614 and the output of the third input buffer 408 are coupled to the output of lighting module 54 by way of multiplexer 1616.

Next provided is a first register unit 1618 having an input coupled to the output of ALU 1614 and an output coupled to the first input of ALU 1614. A second register unit 1620 has an input coupled to the output of ALU 1614. Also, such second register 1620 has an output coupled to the first input and the second input of MLU 1610.

A lighting logic unit (LLU) 1622 is also provided having a first input coupled to the output of ALU 1614, a second input coupled to the output of the first input buffer 404, and an output coupled to the first input of MLU 1610. It should be noted that the second input of LLU 1622 is coupled to the output of the first input buffer 404 via a delay 1624. Further, the output of LLU 1622 is coupled to the first input of MLU 1610 via a first-in first-out register unit 1626. As shown in FIG. 16, the output of LLU 1622 is also coupled to the first input of MLU 1610 via a conversion module 1628. In operation, such conversion module 1628 is adapted for converting scalar vertex data to vector vertex data in a manner similar to that of transform module 52.

Finally, memory 1500 is coupled to at least one of the inputs of MLU 1610 and the output of arithmetic logic unit 1614. In particular, memory 1610 has a read terminal coupled to the first and the second input of MLU 1610. Further, memory 1500 has a write terminal coupled to the output of ALU 1614.

The memory has stored therein a plurality of constants and variables for being used in conjunction with input buffers 402, MLU 1610, ALU 1614, first register unit 1618, second register unit 1620, and LLU 1622 for processing the vertex data.

FIG. 17 is a schematic diagram of MLU 1610 of lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. As shown, MLU 1610 of lighting module 54 includes three multipliers 1700 in parallel. In operation, the present MLU 1610 is adapted to multiply two three component vectors, or pass one three component vector. The multiplication of the three component vectors might be accomplished by way of a dot product or a parallel multiply. Table 13 illustrates the operations that MLU 1610 of lighting module 54 is capable of performing.



TABLE 13

|           |  |
|-----------|--|
| ZMLU_MULT | $o[0] = a[0]*b[0], o[1] = a[1]*b[1], o[2] = a[2]*b[2]$ |
| ZMLU_PASA | $o[0] = a[0], o[1] = a[1], o[2] = a[2]$                |
| ZMLU_PASB | $o[0] = b[0], o[1] = b[1], o[2] = b[2]$                |

Table 14 illustrates the possible A and B inputs of MLU 1610 of lighting module 54.

TABLE 14

|      |                                   |
|------|-----------------------------------|
| MA_V | VBUFFER                           |
| MA_L | LLU                               |
| MA_R | RLU[2,3] (shared with MB_R)       |
| MA_C | Context memory (shared with MB_C) |
| MB_M | MLU                               |
| MB_W | WBUFFER                           |
| MB_R | RLU[2,3] (shared with MA_R)       |
| MB_C | Context memory (shared with MA_C) |

FIG. 18 is a schematic diagram of ALU 1614 of lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. As shown, ALU 1614 includes three adders 1800 in parallel/series. In use ALU 1614 is capable of adding two three component vectors, or passing one three component vector. Table 15 illustrates the various operations of which ALU 1614 of lighting module 54 is capable.

TABLE 15

|            |  |
|------------|--|
| ZALU_ADD   | $o[0] = a[0]+b[0], o[1] = a[1]+b[1], o[2] = a[2]+b[2]$ |
| ZALU_SUM3B | $o[012] = b[0] + b[1] + b[2]$                          |
| ZALU_PASA  | $o[0] = a[0], o[1] = a[1], o[2] = a[2]$                |
| ZALU_PASB  | $o[0] = b[0], o[1] = b[1], o[2] = b[2]$                |

Table 16 illustrates the possible A and B inputs to ALU 1614 of lighting module 54.

TABLE 16

|      |          |
|------|----------|
| AA_W | WBUFFER  |
| AA_R | RLU[0,1] |
| AB_M | MLU      |

FIG. 19 is a schematic diagram of register units 1618 and 1620 of lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. As shown, register units 1618 and 1620 each include two sets of registers 1900 each having an output connected to a first input of a corresponding multiplexer 1902 and an input coupled to a second input of multiplexer 1902.

Register units 1618 and 1620 of lighting module 54 are split into two registers for ALU 1614 and two registers for MLU 1610. In one embodiment, the registers are threaded. The register units 1618 and 1620 exhibit zero latency when a write address is the same as a read address due to a bypass path from the input to the outputs.

FIG. 20 is a schematic diagram of LLU 1622 of lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. LLU 1622 is the lighting unit of lighting module 54. It is a scalar block that computes lighting coefficients later used to multiply the light+material colors. LLU 1622 includes two MAC's, an inverter, four small memories, and a flag register.

The flag register is used to implement the conditional parts of the lighting equations. The outputs are an ambient, diffuse, and specular coefficient. The scalar memories contain variables used for the specular approximations and

constants. The first location of each memory contains 1.0 (for ctx0 and ctx2) and 0.0 (for ctx1 and ctx3). In one embodiment, these are hardwired and do not need to be loaded.

In use LLU 1622 fundamentally implements the equation:  $(x+L)/(M*x+N)$ . This equation is used to approximate a specular lighting term. The inputs to LLU 1622 are from ALU 1614 of lighting module 54 and are the dot products used in the lighting equations. As set forth earlier, with respect to FIG. 16, there is an output FIFO 1626 between LLU 1622 and MLU 1610 which buffers coefficients until MLU 1610 needs them. In one embodiment, such FIFO 1626 might be threaded along with delays 1608 and 1624, and registers 1618 and 1620. Due to possible color material processing, it is unknown when the diffuse and specular outputs are consumed by MLU 1610.

There is specially adapted hardware for dealing with the diffuse output alpha component since lighting module 54 only deals with R,G,B components. Such specially adapted hardware is capable of outputting two types of alpha components, namely vtx color  $\alpha$ [Tbuffer], and stored ctx  $\alpha$ [Ctx store]. The choice between the foregoing alpha components is governed by mode bits 202.

In operation, LLU 1622 calculates ambient (Ca), diffuse (Cde), and specular (Cs) coefficients of lighting. These coefficients are then multiplied with the ambient, diffuse, and specular colors to generate a light's contribution to the vertex color. Table 16A includes a list of inputs received by LLU 1622 and the calculations carried out to generate the ambient (Ca), diffuse (Cde), and specular (Cs) coefficients of lighting. It should be noted that any desired hardware configuration might be employed to implement LLU 1622. In one embodiment, the specific configuration shown in FIG. 20 might be employed.

TABLE 16A

|   |   |
|---|---|
| Input definitions:  |   |
| n = normal vector   | (from transform engine)                 |
| e = normalized eye vector   | (from transform engine)                 |
| l = normalized light vector   | (from transform engine)                 |
| s = spotlight vector*light vector   | (from transform engine)                 |
| D = distance vector (1,d,d*d)   | (from transform engine)                 |
| h = half angle vector   | (from lighting engine)                  |
| K = attenuation constant vector (K0,K1,K2)  | (from context memory)                   |
| The LLU might receive the following scalar data in carrying out its calculations: |   |
| n*1   | (from MLU/ALU)                          |
| n*h   | (from MLU/ALU)                          |
| K*D   | (from MLU/ALU)                          |
| s   | (from transform engine)                 |
| power0  | (material exponent from ctx0-3 memory)  |
| power1  | (spotlight exponent from ctx0-3 memory) |
| range   | (from ctx0-3 memory)                    |
| cutoff  | (from ctx0-3 memory)                    |
| <u>Infinite Light</u>   |   |
| LLU Calculations:   |   |
| Ca = 1.0  |   |
| Cd = n*1  |   |
| Cs = (n*h) <sup>power0</sup>  |   |
| <u>Local Light</u>  |   |
| LLU Calculations:   |   |
| att = 1.0/(K*D)   |   |
| Ca = att  |   |
| Cd = att*(n*1)  |   |
| Cs = att*((n*h) <sup>power()</sup> )  |   |



TABLE 16A-continued

|                                 |
|---------------------------------|
| Spot Light                      |
| LLU Calculations:               |
| $att = (s^{power1}) / (K * D)$  |
| $Ca = att$                      |
| $Cd = att * (n * 1)$            |
| $Cs = att * ((n * h)^{power0})$ |

As set forth above, the mode bits controlling the vertex sequencer might not necessarily be changed by the vertex data itself or by any results derived from vertex data. To allow vertex data to modify vertex processing, LLU 1622 employs a flag register 1623 is provided. Setting bits to TRUE in this flag register allows clamping to 0.0 of calculation results if a flag is specified in the output control of the calculation. Another use of the flag register 1623 would be in setting a write mask for register writes.

The flag register 1623 is provided in LLU 1622 for performing the if/then/else clamping to 0.0 in the lighting equations at no performance penalty. The sign bit of various operands might set the flags. Table 16B illustrates the manner in which the flags in flag register 1623 are set and the resulting clamping.

TABLE 16B

|  |
|--|
| Infinite Light                                   |
| LLU Calculations:                                |
| Dflag = sign bit of (n*1)                        |
| Sflag = sign bit of (n*h)                        |
| Clamp:   |
| $Ca = (0 \quad ) ? 0 : Ca;$                      |
| $Cd = (Dflag \quad ) ? 0 : Cd;$                  |
| $Cs = (Dflag   Sflag) ? 0 : Cs;$                 |
| Local Light                                      |
| LLU Calculations:                                |
| Rflag = sign bit of (range-d)                    |
| Dflag = sign bit of (n*1)                        |
| Sflag = sign bit of (n*h)                        |
| Clamp:   |
| $Ca = (Rflag \quad ) ? 0 : Ca;$                  |
| $Cd = (Rflag   Dflag \quad ) ? 0 : Cd;$          |
| $Cs = (Rflag   Dflag   Sflag) ? 0 : Cs;$         |
| Spot Light                                       |
| LLU Calculations:                                |
| Cflag = sign bit of (s-cutoff)                   |
| Rflag = sign bit of (range-d)                    |
| Dflag = sign bit of (n*1)                        |
| Sflag = sign bit of (n*h)                        |
| Clamp:   |
| $Ca = (Cflag   Rflag \quad ) ? 0 : Ca;$          |
| $Cd = (Cflag   Rflag   Dflag \quad ) ? 0 : Cd;$  |
| $Cs = (Cflag   Rflag   Dflag   Sflag) ? 0 : Cs;$ |

FIG. 21 is an illustration of the organization of the flag register 1623 associated with lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. The flag register 1623 contains 8 one bit flags and are set by the sign bit of the ALU (IFLAG) or MACO (MFLAG) outputs.

When LLU 1622 outputs a scalar value to MLU 1610 where it gets smeared into a tri-word, it specifies a mask for the flag register. If the register & mask is true, 0.0 replaces the output. Table 17 illustrates the various flags of FIG. 21 to be used in outputting ambient, diffuse, and specular attributes.

TABLE 17

|                |              |
|----------------|--------------|
| Ambient Mask:  | C,R, U       |
| Diffuse Mask:  | D, C,R, U    |
| Specular Mask: | D,S, C,R,T,U |

The approximation used for the specular term can go negative where the actual  $\cos(\theta)^n$  would go to 0.0. As a result, it is necessary to perform a clamping operation. For this, the T, U flags are used. Table 18 illustrates various operations of which a functional logic unit (FLU) 1621 of LLU 1622 is capable. Note FIG. 20.

TABLE 18

|            |                           |                               |
|------------|---------------------------|-------------------------------|
| ZFLU_INV   | $o = 1/a$                 | (mantissa accuracy - 12 bits) |
| ZFLU_ISQ   | $o = 1/\sqrt{a}$          | (mantissa accuracy - 6 bits)  |
| ZFLU_PASS  | $o = a$                   |                               |
| ZFLU_PASS1 | $o = 1.0$                 |                               |
| ZFLU_MIN1  | $o = (a < 1.0) ? a : 1.0$ |                               |
| ZFLU_NOP   | $o = 0.0$                 |                               |

FIG. 22 is an illustration of the micro-code fields associated with lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. As shown, the micro-code of lighting module 54 is arranged into 33 fields making up a total width of 85 bits. Fields are delayed to match the data flow of the units. The MLU operations are done at a delay of zero, ALU operations are done at a delay of one, and RLU, LLU output operations are done at a delay of two. Each delay is equivalent to three cycles.

FIG. 23 is a schematic diagram of sequencer 2300 associated with lighting module 54 of FIG. 16 in accordance with one embodiment of the present invention. As shown, sequencer 2300 of lighting module 54 includes an input buffer 2302 adapted for receiving mode bits 202 which are indicative of the status of a plurality of modes of process operations. Also included is memory 1502 capable of storing code segments that each are adapted to carry out the process operations in accordance with the status of the modes.

A sequencing module 2306 is coupled between memory 1502 and buffer 2302 for identifying a plurality of addresses in memory 1502 based on a control vector 2305 derived from the mode bits. The sequencing module 2306 is further adapted for accessing the addresses in memory 1502 for retrieving the code segments that might be used to operate lighting module 54.

The sequencer 2300 of lighting module 54 is similar to that of transform module 52. In operation, sequencer 2300 of lighting module 54 steps through a threaded control vector that is derived from threaded mode bits 202 and executes every ROM address whose corresponding control vector bit is set to "1". The control vector has the same number of bits as the ROM has words. The sequencer 2300 can step through an arbitrary control vector at the rate of a single "1" or enabled bit per a predetermined number of cycles for every thread. Commands that do not use mode bits 202 are executed by on-the-fly micro-code generation. The main difference between sequencer 2300 of lighting module 54 and sequencer 1200 of transform module 52 is that sequencer 2300 of lighting module 54 can loop back and execute the lighting code up to eight times.

The sequencer 2300 of lighting module 54 has a light counter that starts at zero for each new vertex and increments by one at the end of the micro-code sequence. If the LIS field of mode bits 202 contains a "1" in the matching bit field, sequencer 2300 goes back and starts over at the beginning of the lighting micro-code block. This continues



until a zero is found in the LIS field or eight lights have been done. Color accumulation is done by incrementing (per light) the ALU registers that store the diffuse and specular color. Automatic memory address indexing is done using the light counter to fetch the correct parameters for each light.

FIG. 24 is a flowchart delineating the method by which the sequencers of the transform and lighting modules 52 and 54 are capable of controlling the input and output of the associated buffers in accordance with one embodiment of the present invention. As shown, vertex data is initially received in a buffer of a first set of buffers in operation 2420. The buffer in which the vertex data is received is based on a round robin sequence.

Subsequently, in operation 2422, an empty buffer of a second set of buffers is identified also based on a round robin sequence. The transform module 52 is coupled between the first set of buffers and the second set of buffers. When the empty buffer of the second set of buffers is identified, the vertex data is processed in transform module and outputted from transform module to the identified empty buffer of the second set of buffers. Note operations 2424 and 2426.

Similarly, an empty buffer of a third set of buffers, or slots or spaces in memory, are identified based on a round robin sequence in operation 2428. The lighting module 54 is coupled between the second set of buffers and the third set of buffers. When the empty buffer of the third set of buffers is identified, the vertex data is processed in the lighting module, as indicated in operation 2430. The vertex data is subsequently outputted from lighting module 52 to the identified empty buffer of the third set of buffers. See operation 2432. It should be noted that the number of buffers, or slots in memory, is flexible and might be changed.

FIG. 25 is a diagram illustrating the method by which the sequencers of the transform and lighting modules 52 and 54 are capable of controlling the input and output of the associated buffers in accordance with the method of FIG. 24. As shown, the first set of buffers, or input buffers 400, feed transform module 52 which in turn feed the second set of buffers, or intermediate buffers 404, 406. The second set of buffers 404, 406 feed lighting module 54 that drains to memory 2550.

In order carry out the method set forth in FIG. 25, the slots of memory 2550 and the buffers of the first and second set are each assigned a unique identifier upon initially receiving vertex.data. Further, a current state of each buffer is tracked. Such state might include an allocated state, a valid state, an active state, or a done state.

The allocated state indicates that a buffer/slot is already allocated to receive an output of the previous graphics-processing module, i.e. transform module or lighting module. When a write pointer is scanning the buffers/slots in the round robin sequence, a buffer/slot in the allocated state cause such write pointer to increment to the next buffer or slot.

If a buffer/slot is in the valid state, the buffer/slot is available for receiving vertex data. On the other hand, the active state indicates that a buffer/slot is currently in an execution state, or receiving vertex data. This active status is maintained until a thread is done after which a read pointer increments, thus placing the buffer/slot back in the valid state. It should be noted that the first set of buffers 400 are only capable of being in the valid state since there is no previous graphics-processing module to allocate them.

An example of a sequence of states will now be set forth. Upon receiving vertex data in one of the first set of buffers 400 and a new set of command bits 200, such buffer is placed in the valid state, after which one of the second set

of buffers 402, 404 is placed in the allocated state in anticipation of the output of transform module 52.

If none of the second set of buffers 404, 406 is available for allocation, the vertex data in the buffer of the first set 400 can not be processed. Further, a check might be done to determine whether the code segments to be executed will interfere with any other code segments that are to be simultaneously run. If so, the vertex data in the buffer of the first set 400 will not be processed and a stall condition initiated.

After one of the second set of buffers 404, 406 is placed in the allocated state, the buffer of the first set 400 is placed in the active state. When transform module 52 is finished execution, the buffer of the second set 404, 406 is read and then placed in the valid state. These state changes are similarly executed during the transfer of vertex data between the second set 404, 406 and the slots of memory 2550.

FIG. 25B illustrates the rasterizer module 56 that comprises a set-up module 57 and a traversal module 58. The rasterizer module 56 is adapted for performing area-based rasterization in an alternating manner. In particular, a plurality of polygon-defining sense points are positioned on or near the primitive after which line equations are evaluated at the points to determine which pixels reside in the primitive. During operation, this evaluation is repeated as the points are moved in an alternating manner for efficiency purposes. Further, the rasterizer module 56 might be adapted to operate without any clipping procedure.

FIG. 26 illustrates a schematic of the set-up module 57 of rasterization module 56. As shown, the set-up module 57 includes a control section 61 that handles routing data and control signals to their appropriate functional units in order to perform the desired floating-point calculations. The primitive sequencer 62 handles turning sequences of vertices into triangles, lines or points. Further, floating point data path section 64 includes the multiplexers and floating point computation units that perform the math required in the set-up unit.

With continuing reference to FIG. 26, output formatting section 63 handles converting the internal floating point format of edge slopes and edge values into integer formats suitable for the rasterizer since the rasterizer operates only with integer values. Of course, in alternate embodiments, the rasterizer might use a floating point thus obviating the need for output formatting section 63.

In operation, output formatting section 63 executes a block floating point conversion. As is well known, with a given number, i.e.  $2.34 \times 10^{10}$ , floating point format tracks a mantissa (2.34) and an exponent (10) thereof. Block floating point conversion essentially manipulates the decimal place of the mantissas of incoming data such that the exponents are the same. To this end, the exponent need not be handled in rasterizer module 56.

FIG. 26A is an illustration showing the various parameters calculated by set-up module 57 of rasterizer module 56 of FIG. 25B. Such parameters are required for rasterizer module 56 to perform the associated functions. Upon receipt of a primitive 2600, set-up module 57 calculates three values including slopes 2601 of the primitive 2600, a starting position 2602 and a starting value 2604.

The slopes 2601 are used to generate coefficients for line equations of the edges of the primitive 2600 to be used during rasterization. The slopes 2601 might, for example, be calculated by using equations #4 and #5 shown below.

$$\text{slope}_A = y_0 - y_1$$

$$\text{slope}_B = x_1 - x_0$$

Equations #4 and #5



where  $y_0, y_1$  and  $x_0, x_1$  are coordinates of vertices shown in FIG. 26A.

It should be noted that the slopes might also be calculated using the coordinates of the vertices by employing a simple rotation operation or the like.

The starting position **2602** indicates a starting point for area rasterization that will be set forth hereinafter in greater detail. The starting value **2604** is equal to the area of the shaded triangle shown in FIG. 26A and is also used during the area-based rasterization process. Such starting value **2604** is selected so that stepping the raster position about the screen while adding the slope at each step will equal zero exactly when the raster position is on the edge. Calculation of the starting value **2604** might be accomplished using Equation #6 below:

$$\text{starting\_value} = \text{slope}_A * (x_s - x_0) + \text{slope}_B * (y_s - y_0) \quad \text{Equation \#6}$$

where

$x_s, y_s$  starting position **2602**

$\text{slope}_A, \text{slope}_B$  = slopes of one of the edges based on coordinates of vertices shown in FIG. 26A

$x_0, y_0$  = coordinates of one of the vertices of the edges shown in FIG. 26A

It should be understood that the foregoing values might also be calculated for other types of primitives. For example, in the case of a line, an extra slope must be calculated for the four-sided bounding box. Such slope can be easily calculated by taking the reciprocal of the slope of an opposite side of the bounding box. In addition to the extra slope calculation, it is noted that another starting value needs to be calculated in the case of the line primitive.

FIG. 27 illustrates the method by which rasterizer module **56** handles one of a plurality of primitives, e.g. triangles. In particular, an initial operation is first performed by set-up module **57** of rasterizer module **56**. Upon receipt of a primitive, line equation coefficients of line equations are determined for lines that define the primitive in operation **2700** using slopes **2601** of FIG. 26A in a manner that is well known to those with ordinary skill in the art. As is well known, three line equations are required to define a triangle. On the other hand, a primitive such as a line is drawn as a rectangle or parallelogram with four sides and four line equations.

Thereafter, in operation **2702**, the line equation coefficients are modified if any primitive vertex(es) has a negative W-coordinate. Additional information regarding this process will be set forth hereinafter in greater detail with reference to FIG. 32.

It should be noted that set-up module **57** of rasterizer module **56** also computes a bounding box of the primitive. For most triangles, the bounding box includes the minimum and maximum values of the three vertexes. For lines, the four parallelogram corners of the bounding box are calculated. For triangles or lines that have a vertex with a negative W-coordinate, an area that is to be drawn extends beyond the convex hull of the vertices.

One of the commands of OpenGL® is a scissor rectangle which defines a boundary outside of which is not to be drawn. The set-up module **57** of rasterizer module **56** calculates the intersection of the bounding box and the scissor rectangle. Since the scissor rectangle is a rectangle, four additional line equations are afforded. It should be noted that the line equations associated with the scissor rectangle have a trivial form, i.e. horizontal or vertical.

Furthermore, in 3-D space, the near plane and far plane are parallel and at right angles to the line of sight. In the case

of the primitive being a triangle, three vertexes are included which define a plane that might have any orientation. The intersections of the plane of the primitive and the near and far planes include two lines with two associated line equations.

Accordingly, each primitive has a total of nine or ten line equations depending on whether it takes the form of a triangle or a line, respectively. Again, in the case of the triangle, such line equations include the three line equations which define the triangle, the four line equations defining the bounding box and the two line equations which define the intersections of the plane in which the primitive resides, and near and far planes.

With continuing reference to FIG. 27, the process progresses in operation **2704** by positioning a plurality of points on or near the primitive. The starting position **2602** dictates such positioning, as shown in FIG. 26A. Such points define an enclosed convex region and reside at corners of the convex region. FIG. 27A illustrates such sense points **2705** that enclose convex region **2707**, e.g. a rectangle. In one embodiment, such rectangle might be 8x2 pixels in size. Further, the points might be initially positioned to enclose a top vertex of the primitive. As an option, this might be accomplished using truncation.

Once the primitive is positioned, the process is continued by traversal module **58** which begins in operation **2706** by processing rows of the primitive in a manner set forth below. After the processing of each row, it is determined whether a jump position has been found in decision **2708**. A jump position is a starting position for processing the next row and will be described hereinafter in greater detail. If it is determined in decision **2708** that a jump position has been found, the sense points that define the convex region are moved thereto in operation **2710**. If, however, it is determined that a jump position has not been found, the process is ended. It should be noted that, in an alternate embodiment, columns, diagonals or any other type of string might be processed in operation **2706** instead of rows.

FIG. 28 is a flowchart illustrating a process of the present invention associated with the process row operation **2706** of FIG. 27. As shown, the process begins by computing the sense points in operation **2800** in order to determine whether the polygon-defining sense points might be moved right in decision **2801**. Such decision is made based on the position of the rightmost sense points. If the rightmost sense points are not positioned outside the same edge or edges of the primitive, rightward movement is permitted and a position (X and Y coordinates) to the right of the current position is stored as a snap location in operation **2802**. If, however, both rightmost sense points are positioned outside one or more edges of the primitive, rightward movement is not permitted and operation **2802** is skipped.

Next, the line equations are evaluated at the points of the convex region, e.g. rectangle, in operation **2804**. The evaluation includes determining if the points reside in the primitive. Such determination as to whether the points reside in the primitive might include determining whether the evaluation of each of the line equations renders a positive value or a negative value at each of the sense points.

The line equations can be formulated to be positive inside the primitive and negative outside. Inclusive edges, for which pixels that lie exactly on the edge should be drawn, evaluate to zero and might be treated as positive. Exclusive edges, which should not be drawn, can be made negative by initially subtracting a value of one from the starting line equation value. Thus pixels on exclusive edges evaluate to a negative value (-1) instead of a positive zero. This permits



the sense point interpretation to ignore the inclusive/exclusive policy and just test the line equation sign.

After the line equations are evaluated at the points, it is determined whether a current position of the sense points constitutes a jump position in decision **2806**. It should be noted that a jump position is stored only if the two bottom sense points are not both outside an edge. If it is determined in decision **2806** that a jump position has been found, the jump position is calculated and stored (or replaces a previously stored jump position if existent) in operation **2808**. If not, however, operation **2808** is skipped.

With continuing reference to FIG. **28**, it is then determined in decision **2810** whether leftmost sense points are both outside an edge of the primitive. Again, this process entails determining whether the evaluation of the line equations at both of the leftmost sense points renders positive or negative values. In particular, upon computation of the coefficients of the nine or ten edge equations at the pertinent sense points, nine or ten values are rendered that have nine or ten sign bits. To determine if the current side is completely outside any edge, for example, the present invention AND's the ten sign bits from the two sense points together. If any bit(s) survive, then both points are outside that edge.

If it is determined that the leftmost sense points are not both outside an edge of the primitive, it is concluded that there still remains further portions of the primitive to be considered in the leftward direction, and the sense points are moved left in operation **2812**. If it is determined in decision **2810** that both leftmost sense points are indeed outside the edge of the primitive, it is concluded that there no longer remains further portions of the primitive to be considered in the leftward direction. Next, in decision **2814**, it is determined whether there is a snap location that resulted from operation **2802**.

If it is determined in decision **2814** that a snap location does not exist, the process is done. If, however, a snap location does exist, the sense points are moved to the snap location in operation **2816**. Thereafter, operations similar to those of operations **2804–2812** are executed to map a right side of the primitive. This begins in operation **2818** by the line equations being evaluated at the points of the convex region.

After the line equations are evaluated at the points, it is determined whether a current position of the sense points constitutes a jump position in decision **2820**. If it is determined in decision **2806** that a jump position has been found, the jump position is calculated and stored in operation **2822**. If not, however, operation **2822** is skipped.

With continuing reference to FIG. **28**, it is then determined in decision **2824** whether rightmost sense points are both outside an edge of the primitive. If it is determined that the rightmost sense points are not both outside an edge of the primitive, it is concluded that there still remains further portions of the primitive in the rightward direction to be considered, and the sense points are moved right in operation **2826**. If it is determined in decision **2824** that both rightmost sense points are outside the edge of the primitive, it is concluded that there no longer remains further portions of the primitive to be considered in the rightward direction, and the instant process is done.

FIGS. **28A** and **28B** are illustrations of the sequence in which the sense points of the present invention might be moved about the primitive **2850**. It should be noted that various alterations might include determining whether the points can go left in decision **2800** and proceeding right initially. Further, the line equations might be defined to indicate whether the points are inside or outside the primitive in any desired way.

To avoid stepping in a repeating loop, the present invention thus employs an overall direction of movement during rasterization. The initial implementation proceeds top-down, visiting every convex region on a row before stepping down to the next. By processing rows top-down as well as never stepping right then left or left then right, loops are thus avoided.

An example of the foregoing process might be shown with reference to the polygon-defining points, **P1**, **P2**, **P3** and **P4** of FIG. **27A**. In operation, pairs of adjacent sense points can be examined to determine whether stepping in their direction would be productive. For example, if both **P3** and **P4** in FIG. **27A** were outside of an edge of a polygon, but **P1** and/or **P2** are not, then clearly the drawable inside region lies to the left, not to the right. Thus the sense points should not move right. Conversely, if both **P3** and **P4** are inside all the edges, then there is a drawable area just beyond **P3** and **P4**, and stepping right is appropriate. Indeed, if **P3** and **P4** were not outside the same edge or edges, stepping right would be productive. This same logic applies to stepping upwards guided by **P1** and **P3**, or stepping left guided by **P1** and **P2**, or stepping downwards based on **P2** and **P4**.

The foregoing process thus moves, or steps, the convex region defined by the points around the inside of the primitive, using sense points as a guide. Since the convex region defined by the points might be large, many pixels might be tested simultaneously. During use, if all sense points are inside all edges of the primitive, then all the enclosed pixels must be drawable (assuming a convex primitive). A significant advantage is afforded by testing the corners, namely the ability of proving an arbitrary area of the primitive is inside, outside or split. Only in the latter case do the individual pixels in the convex region defined by the points need to be tested. In such case, the pixels in the convex region defined by the points might be tested one-by-one or by another method in order to determine whether they reside in the primitive. Furthermore, the sense points might reduce the amount of further testing required by defining which edges(s) split the area and which do not.

FIG. **29** is a flowchart illustrating an alternate boustrophedonic process of the present invention associated with the process row operation **2706** of FIG. **27**. As shown, it is first determined in decision **2900** whether a previous movement was in a first or second direction. If there was not any actual previous movement, a default previous movement might be assumed. If it is determined in decision **2900** that the previous movement was in a second direction, the line equations are evaluated at the points of the convex region, e.g. a rectangle, in operation **2902** in a manner similar to operation **2804** of FIG. **28**.

With continuing reference to FIG. **29**, it is subsequently determined in decision **2904** as to whether sense points of a first side of the rectangle are both outside an edge of the primitive. If not, the sense points are moved or stepped in the first direction in operation **2906**. Upon it being determined that the sense points of the first side of the rectangle are both outside an edge of the primitive, it is then determined in decision **2905** whether the points can be moved downwardly or, in other words, whether the current position constitutes a jump position. If so, a jump position is calculated and stored in operation **2908** after which the process is done.

On the other hand, if it is determined in decision **2900** that the previous movement was in a first direction, operations similar to those of operation **2902–2908** are carried out. In particular, the line equations are evaluated at the points of the convex region, e.g. a rectangle, in operation **2910**. It is then determined in decision **2912** as to whether sense points



of a second side of the rectangle are both outside an edge of the primitive. If not, the sense points are moved or stepped in the second direction in operation 2914. Upon it being determined that the sense points of the second side of the rectangle are both outside an edge of the primitive, it is then determined in decision 2913 whether the points can be moved downwardly or, in other words, whether the current position constitutes a jump position. If so, a jump position is calculated and stored in operation 2916 after which the process is done.

FIG. 29A is an illustration of the sequence in which the sense points of the present invention are moved about the primitive in accordance with the boustrophedonic process of FIG. 29. The foregoing boustrophedonic rasterization constrains the sequence to obey certain rules that offer better performance for hardware. As shown, the boustrophedonic rasterization affords a serpentine pattern that folds back and forth. A horizontal boustrophedonic sequence, for example, might generate all the pixels within a primitive triangle that are on one row from left to right, and then generate the next row right to left, and so on. Such a folded path ensures that an average distance from a generated pixel to recently previously generated pixels is relatively small.

Generating pixels that are near recently previously generated pixels is important when recent groups of pixels and/or their corresponding texture values are kept in memories of a limited size. The boustrophedonic sequence more often finds the pixels or texture values already loaded into such memories, and therefore repeating the memory load occurs less often.

As an option, at least one boundary might be used which divides the primitive into a plurality of portions prior to rasterization. In operation, the points might be moved in each of the portions separately. Further, the points might be moved through an entirety of a first one of the portions before being moved in a second one of the portions.

FIG. 30 is a flowchart illustrating an alternate boustrophedonic process using boundaries. As an option, the decision whether to use boundaries might be based on a size of the primitive. As shown in FIG. 30, the boustrophedonic process which handles boundaries is similar to that of FIG. 27 with the exception of an additional operation 3000 wherein at least one boundary is defined which divides the primitive into a plurality of portions or swaths.

With continuing reference to FIG. 30, an additional decision 3001 follows the completion of every portion of the primitive. In particular, it is determined in decision 3001 whether a start position of an adjacent portion was found in operation 3006. If so, the convex region defined by the sense points is moved to a start position of an adjacent portion of the primitive in operation 3002 and operations 3004–3010 are repeated for the new portion of the primitive. Further information relating to the determination of the start position in operation 3006 will be set forth in greater detail during reference to FIG. 31.

FIG. 31A is an illustration of the process by which the convex region of the present invention is moved about the primitive in accordance with the boundary-based boustrophedonic process of FIG. 30. As shown, the first portion that is processed is that which includes the topmost vertex of the primitive. During operation, a left neighboring portion is processed after which the adjacent left neighboring portion is processed and so on. This is continued until there are no remaining left neighboring portions. Next, a neighboring portion to the right of the first portion is processed after which the adjacent right neighboring portion is processed and so on until all of the right neighboring portions are

processed. It should be appreciated that other types of ordering schemes might be utilized per the desires of the user.

FIG. 31 is a flowchart showing the process associated with the process row operation 3006 of FIG. 30. Such process is similar to the boustrophedonic process of FIG. 29 with the exception of decisions 3118 through 3121. Decisions 3118 and 3120 both determine whether any of the sense points have passed any boundary. Only if it is determined that the sense points are still within the boundaries is the respective loop continued.

In operations 3119 and 3121, starting positions of adjacent portions of the primitive are sought and stored when it is determined in decisions 3118 and 3120 that any sense points of the convex region have passed any boundary, respectively. As shown in FIG. 31A, such starting positions 3126 are each defined as being the topmost point of a portion of the primitive existent beyond a boundary. By storing this position, a starting point is provided when the process is repeated for the adjacent boundary-defined portion of the primitive.

It should be noted that operations 3119 and 3121 are both performed while processing the first portion of the primitive. While not expressly shown in FIG. 31, only a first one of such operations is performed when processing portions to the left of the first portion, while only a second one of such operation is performed when processing portions to the right of the first portion. In other words, when processing portions to the left of the first portion, starting positions are only determined when a leftmost boundary of the currently processed portion has been exceeded. Similarly, when processing portions to the right of the first portion, starting positions are only determined when a rightmost boundary of the currently processed portion has been exceeded.

Using boundaries during rasterization solves a very critical problem during pipeline processing. If a primitive is very wide, the storage associated with the pixels of a single row might not fit in a limited-size memory. Rasterization with boundaries divides the triangle into limited-width rows (or columns), and generates all the pixels within such a portion before moving on to the next portion.

For example, even if a triangle is 100 pixels wide, a limited-size pixel or texture memory might only hold information for the previous 20 pixels. Constraining the pixel sequence to stay within ten-pixel-wide vertical portions allows all the pixels on the previous and current rows to fit in the memory. This means that a boustrophedonic sequence within a boundary-defined portion would always have the previous pixel on the current row (if any) in the memory, as well as the pixels in the row above (if any) in the memory as well.

Most underlying memory systems transfer blocks of data with a certain overhead per block. Small accesses to the memory system are penalized heavily by this overhead. In order to be efficient, larger accesses are employed and the rest of the block is maintained in case it might be used next. Beyond that, a cache memory system keeps a plurality of these recent blocks, increasing the probability that memory accesses can be avoided.

The boustrophedonic sequence of the present invention exploits the single-retained-block concept when it reverses and handles pixels immediately below one end of the current line. Further, the boustrophedonic sequence exploits cache when it limits rasterization to portions of a particular size. Specifically, two scanlines within a portion should fit in the cache, so throughout the second scanline, benefits might be incurred from cache storage of the first scanline.



There is no constraint on the sequence or number of boundary-defined portions. Although the present description uses the example of vertical portions and a horizontal boustrophedonic pattern, similar principles might extend to horizontal portions, vertical boustrophedonic patterns or even to diagonal portions and patterns. In one embodiment, the length of the strings (e.g. rows, columns, diagonals, etc.) might be each limited to be less than a dimension of the primitive along which the string resides.

FIG. 32 is a flowchart showing the process associated with operation 2702 of FIG. 27. The instant process is designed to handle a primitive with portions that reside behind the eye. These outlying portions might cause problems in subsequent rasterization operations. To accomplish this, the instant process employs a variable,  $W$  that is commonly used for projection i.e., for viewing objects in perspective. The variable  $W$  is a number that the other coordinates,  $X$ ,  $Y$  and  $Z$ , are divided by in order to make nearby things larger and far things smaller. The variable  $W$  is representative of a distance between a center of projection and the corresponding vertex.

As shown in FIG. 32, a primitive is first received that is defined by a plurality of vertices. Each of such vertices includes a  $W$ -value. Upon the receipt of the primitive, the set-up module serves to define lines that characterize the primitive based on the vertices. Note operation 3200.

The  $W$ -values are then analyzed in decision 3202. As shown, if one of the  $W$ -values is negative, a line equation for a line opposite the vertex having the negative value is flipped in operation 3204. In other words, the coefficients of the line equation are multiplied by  $-1$ . Further, if two of the  $W$ -values are negative, line equations for lines connecting the vertex having a positive  $W$ -value and each of the vertexes having negative  $W$ -values are flipped in operation 3206. If three of the  $W$ -values are negative, a cull condition 3207 occurs where the present invention culls the triangle. Still yet, if none of the  $W$ -values are negative, no additional action is taken.

FIGS. 32A–32C illustrate the manner in which flipping line equations affects which portion of the screen is processed. FIG. 32A shows the case where none of the  $W$ -values are negative and the line equations are left unaltered. As shown, an interior portion of the primitive is filled in such case.

FIG. 32B shows the case where one of the  $W$ -values is negative and which of the line equations is flipped accordingly. As shown, the portion of the primitive opposite the vertex is filled in the present case. In particular, the area to be drawn is bounded by two lines that are co-linear with the two triangle sides sharing the  $-W$  vertex, and further bounded by a side of the triangle that shares the two  $+W$  vertexes.

FIG. 32C shows the case where two of the  $W$ -values are negative and which of the line equations are flipped accordingly. As shown, the portion of the primitive opposite the vertexes is filled using the methods and/or processes set forth hereinabove with reference to FIGS. 27–32. In other words, the area to be drawn is bounded by two lines that are co-linear with the two triangle sides sharing the  $+W$  vertex, and further contiguous to the  $+W$  vertex.

The present invention is thus capable of handling all three of the foregoing cases. If part of the triangle is beyond the near and/or far plane, it draws only the portion within those planes. If the triangle has one or two negative  $Z$  vertexes, only the correct  $+Z$  portion is drawn.

Even if all vertexes are off-screen, and the triangle extends from behind the eye to beyond the far plane,

whatever pixels are inside the triangle and on the screen and have  $Z$  between the near and far limits. The present invention ensures that little time is wasted exploring bad pixels. This is possible because all clipping, by screen edge or the near or far plane, always results in a convex region on-screen which can be explored easily.

A problem sometimes arises when the starting point is not inside the area to be filled. This can occur if the top vertex is off-screen or is clipped by the near or far plane. In this case, the traversal stage must search for the top point of the drawn region, starting from above. It can do this efficiently by being guided by the signs of the triangle edge slopes and the  $Z$  slope. It can test the triangle line equations to discover it is outside the drawn region and why. When it knows what edge(s) and/or  $Z$  limit it is outside of, it knows what direction(s) to step that brings it closer to that edge or limit. By moving horizontally in preference to vertically (when there is a choice), searching for the drawn region guarantees it finds the top drawable pixel if there is one. This problem also occurs with external ( $-W$ ) triangles that open up. In this case, the drawn area extends above all three vertexes.

In one embodiment of the present invention, traversal proceeds from top to bottom of the triangle. The starting point is the top vertex of the triangle if none have a negative  $W$ -value and the top vertex is in the scissor rectangle. Otherwise, a point on the top of the scissor rectangle is chosen. Since traversal always begins within the scissor rectangle and never ventures out of it, only the portion of the triangle within the scissor rectangle is ever drawn, even if the area enclosed by the edges extends far beyond the scissor rectangle. In this way, simple scissor rectangle-edge clipping is effected.

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A method for performing rasterization using alternating point traversal, comprising:

receiving a primitive;

positioning a plurality of points on or near the primitive, wherein the points define an enclosed convex region;

defining at least one boundary which divides the primitive into a plurality of portions; and

moving the points in an alternating manner in the portions via at least one of a plurality of contiguous rows and a plurality of contiguous columns for the purpose of identifying an area in the primitive.

2. The method as recited in claim 1, wherein the points are moved in a boustrophedonic manner.

3. The method as recited in claim 1, wherein the convex region is a polygon defined by the points.

4. The method as recited in claim 1, wherein the points are also moved based on the evaluation of line equations for lines that define the primitive.

5. The method as recited in claim 1, wherein the points are moved in each of the portions separately.

6. The method as recited in claim 5, wherein the points are moved through essentially an entirety of a first one of the portions before being moved in a second one of the portions.

7. The method as recited in claim 1, wherein a boundary spacing is defined based on a memory size.

8. A computer program embodied on a computer readable medium for performing rasterization using alternating point traversal, comprising:



a code segment for receiving a primitive;  
 a code segment for positioning a plurality of points on or near the primitive, wherein the points define an enclosed convex region;  
 a code segment for defining at least one boundary which divides the primitive into a plurality of portions; and  
 a code segment for moving the points in an alternating manner in the portions via at least one of a plurality of contiguous rows and a plurality of contiguous columns for the purpose of identifying an area in the primitive.

9. The computer program as recited in claim 8, wherein the points are moved in a boustrophedonic manner.

10. The computer program as recited in claim 8, wherein the convex region is a polygon defined by the points.

11. The computer program as recited in claim 8, wherein the points are also moved based on the evaluation of line equations for lines that define the primitive.

12. The computer program as recited in claim 8, wherein the points are moved in each of the portions separately.

13. The computer program as recited in claim 12, wherein the points are moved through essentially an entirety of a first one of the portions before being moved in a second one of the portions.

14. The computer program as recited in claim 8, wherein a boundary spacing is defined based on a memory size.

15. A system for performing rasterization using alternating point traversal, comprising:

(a) a set-up module for receiving a primitive, and positioning a plurality of points on or near the primitive, wherein the points define an enclosed convex region; and

(b) a traversal module coupled to the set-up module for moving the points in an alternating manner for the purpose of identifying an area in the primitive; wherein at least one boundary is defined which divides the primitive into a plurality of portions, and the points are moved in the portions via at least one of a plurality of contiguous rows and a plurality of contiguous columns.

16. The system as recited in claim 15, wherein the points are moved in a boustrophedonic manner.

17. The system as recited in claim 15, wherein the convex region is a polygon defined by the points.

18. The system as recited in claim 15, wherein the points are also moved based on the evaluation of line equations for lines that define the primitive.

19. The system as recited in claim 15, wherein the points are moved in each of the portions separately.

20. The system as recited in claim 19, wherein the points are moved through essentially an entirety of a first one of the portions before being moved in a second one of the portions.

21. The system as recited in claim 15, wherein a boundary spacing is defined based on a memory size.

22. A method for performing rasterization, comprising: receiving a primitive; identifying pixels in the primitive for rendering purposes by stepping along adjacent strings of adjacent pixels; and limiting a length of the strings of pixels which are stepped; wherein the length of the strings is limited by defining at least one boundary which divides the primitive into a plurality of portions each including at least one of a plurality of contiguous rows and a plurality of contiguous columns.

23. The method as recited in claim 22, wherein a plurality of points are moved through essentially an entirety of a first one of the portions before being moved in a second one of the portions.

24. The method as recited in claim 23, and further comprising identifying a starting position of the second portion while moving through the first portion.

25. The method as recited in claim 22, wherein a boundary spacing is defined based on a memory size.

26. The method as recited in claim 22, wherein the length of the strings are each limited to be less than a dimension of the primitive along which the string resides.

27. A system for performing rasterization, comprising:

(a) a set-up module for receiving a primitive; and

(b) a traversal module coupled to the set-up module for identifying pixels in the primitive for rendering purposes by stepping along adjacent strings of adjacent pixels, wherein a length of the strings is limited; wherein the length of the strings is limited by defining at least one boundary which divides the primitive into a plurality of portions each including at least one of a plurality of contiguous rows and a plurality of contiguous columns.

28. The system as recited in claim 27, wherein a plurality of points are moved through essentially an entirety of a first one of the portions before being moved in a second one of the portions.

29. The system as recited in claim 28, and further comprising identifying a starting position of the second portion while moving through the first portion.

30. The system as recited in claim 29, wherein a boundary spacing is defined based on a memory size.

31. The system as recited in claim 27, wherein the length of the strings are each limited to be less than a dimension of the primitive along which the string resides.

32. A method for performing rasterization, comprising: receiving a primitive; identifying pixels in the primitive for rendering purposes by stepping along adjacent strings of adjacent pixels; and limiting a length of the strings of pixels which are stepped; wherein a length of the strings is limited by defining at least one boundary which divides the primitive into a plurality of portions, and a boundary spacing is defined based on a memory size.

33. A system for performing rasterization using alternating point traversal, comprising:

(a) a set-up module for receiving a primitive, and positioning a plurality of points on or near the primitive, wherein the points define an enclosed convex region; and

(b) a traversal module coupled to the set-up module for moving the points in an alternating manner for the purpose of identifying an area in the primitive; wherein a length of the strings is limited by defining at least one boundary which divides the primitive into a plurality of portions, and a boundary spacing is defined based on a memory size.

34. The method as recited in claim 1, wherein the at least one of a plurality of contiguous rows and a plurality of contiguous columns includes a plurality of contiguous rows.

35. The method as recited in claim 1, wherein the at least one of a plurality of contiguous rows and a plurality of contiguous columns includes a plurality of contiguous columns.

UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 6,650,325 B1  
APPLICATION NO. : 09/454505  
DATED : November 18, 2003  
INVENTOR(S) : Voorhies et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the specification:

col. 1, line 19 replace “ans” with --and--;

In the claims:

col. 37, line 36 replace “portions,;” with --portions,--.

Signed and Sealed this

Twenty-third Day of February, 2010

A handwritten signature in black ink that reads "David J. Kappos". The signature is written in a cursive, flowing style.

David J. Kappos  
*Director of the United States Patent and Trademark Office*